

Findings of the Association for Computational Linguistics: ACL 2023

Toronto, Canada
9 – 14 July 2023

Volume 1 of 17

ISBN: 978-1-7138-8241-1

Printed from e-media with permission by:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571



Some format issues inherent in the e-media version may also appear in this print version.

Copyright© (2023) by the Association for Computational Linguistics
All rights reserved.

Printed with permission by Curran Associates, Inc. (2024)

For permission requests, please contact the Association for Computational Linguistics
at the address below.

Association for Computational Linguistics
209 N. Eighth Street
Stroudsburg, Pennsylvania 18360

Phone: 1-570-476-8006

Fax: 1-570-476-0860

acl@aclweb.org

Additional copies of this publication are available from:

Curran Associates, Inc.
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: 845-758-0400
Fax: 845-758-2633
Email: curran@proceedings.com
Web: www.proceedings.com

Table of Contents

<i>Investigating Glyph-Phonetic Information for Chinese Spell Checking: What Works and What's Next?</i> Xiaotian Zhang, Yanjun Zheng, Hang Yan and Xipeng Qiu	1
<i>A Self-Supervised Integration Method of Pretrained Language Models and Word Definitions</i> Hwiyeol Jo	14
<i>Conformal Nucleus Sampling</i> Shauli Ravfogel, Yoav Goldberg and Jacob Goldberger	27
<i>DiscoPrompt: Path Prediction Prompt Tuning for Implicit Discourse Relation Recognition</i> Chunkit Chan, Xin Liu, Jiayang Cheng, Zihan Li, Yangqiu Song, Ginny Y. Wong and Simon See	35
<i>Modularized Zero-shot VQA with Pre-trained Models</i> RUI CAO and Jing Jiang	58
<i>TimelineQA: A Benchmark for Question Answering over Timelines</i> Wang-Chiew Tan, Jane Dwivedi-Yu, Yuliang Li, Lambert Mathias, Marzieh Saeidi, Jing Nathan Yan and Alon Y. Halevy	77
<i>Abstractive Text Summarization Using the BRIO Training Paradigm</i> Khang Nhut Lam, Thieu Gia Doan, Khang Thua Pham and Jugal Kalita	92
<i>Modeling the Q-Diversity in a Min-max Play Game for Robust Optimization</i> Ting Wu, Rui Zheng, Tao Gui, Qi Zhang and Xuanjing Huang	100
<i>Pre-training Language Model as a Multi-perspective Course Learner</i> Beiduo Chen, Shaohan Huang, Zihan Zhang, Wu Guo, Zhenhua Ling, Haizhen Huang, Furu Wei, Weiwei Deng and Qi Zhang	114
<i>Layerwise universal adversarial attack on NLP models</i> Olga Tsymbol, Danil Malaev, Andrei Petrovskii and Ivan Oseledets	129
<i>Scene-robust Natural Language Video Localization via Learning Domain-invariant Representations</i> zehan wang, Yang Zhao, Haifeng Huang, yan xia and Zhou Zhao	144
<i>Exploiting Pseudo Image Captions for Multimodal Summarization</i> Chaoya Jiang, Rui Xie, Wei Ye, Jinan Sun and Shikun Zhang	161
<i>Cross-Lingual Transfer with Target Language-Ready Task Adapters</i> Marinela Parovic, Alan Ansell, Ivan Vulić and Anna Korhonen	176
<i>DynaMiTE: Discovering Explosive Topic Evolutions with User Guidance</i> Nishant Balepur, Shivam Agarwal, Karthik Venkat Ramanan, Susik Yoon, Diyi Yang and Jiawei Han	194
<i>Boost Transformer-based Language Models with GPU-Friendly Sparsity and Quantization</i> Chong Yu, Tao Chen and Zhongxue Gan	218
<i>RMSSinger: Realistic-Music-Score based Singing Voice Synthesis</i> Jinzheng He, Jinglin Liu, Zhenhui Ye, Rongjie Huang, Chenye Cui, Huadai Liu and Zhou Zhao	236

<i>Zero-Shot Prompting for Implicit Intent Prediction and Recommendation with Commonsense Reasoning</i> Hui-Chi Kuo and Yun-Nung Chen	249
<i>MTGP: Multi-turn Target-oriented Dialogue Guided by Generative Global Path with Flexible Turns</i> Anqi Liu, Bo Wang, Yue Tan, Dongming Zhao, Kun Huang, Ruifang He and Yuexian Hou ..	259
<i>The Larger they are, the Harder they Fail: Language Models do not Recognize Identifier Swaps in Python</i> Antonio Valerio Miceli Barone, Fazl Barez, Shay B. Cohen and Ioannis Konstas	272
<i>Class Lifelong Learning for Intent Detection via Structure Consolidation Networks</i> Qingbin Liu, Yanchao Hao, Xiaolong Liu, Bo Li, Dianbo Sui, Shizhu He, Kang Liu, Jun Zhao, Xi Chen, Ningyu Zhang and Jiaoyan Chen	293
<i>On Evaluating and Mitigating Gender Biases in Multilingual Settings</i> Aniket Vashishtha, Kabir Ahuja and Sunayana Sitaram	307
<i>Rethinking Round-Trip Translation for Machine Translation Evaluation</i> Terry Yue Zhuo, Qiongfai Xu, Xuanli He and Trevor Cohn	319
<i>G^oxt3R: A Graph-Guided Generate-and-Rerank Framework for Complex and Cross-domain Text-to-SQL Generation</i> Yanzheng Xiang, Qian-Wen Zhang, Xu Zhang, Zejie Liu, Yunbo Cao and Deyu Zhou	338
<i>A Unified Knowledge Graph Augmentation Service for Boosting Domain-specific NLP Tasks</i> Ruiqing Ding, Xiao Han and Leye Wang	353
<i>Dialogue Planning via Brownian Bridge Stochastic Process for Goal-directed Proactive Dialogue</i> Jian Wang, Dongding Lin and Wenjie Li	370
<i>A Match Made in Heaven: A Multi-task Framework for Hyperbole and Metaphor Detection</i> Naveen Badathala, Abisek Rajakumar Kalarani, Tejpalsingh Siledar and Pushpak Bhattacharyya	388
<i>Prompt Tuning for Unified Multimodal Pretrained Models</i> Hao Yang, Junyang Lin, An Yang, Peng Wang and Chang Zhou	402
<i>Learning Joint Structural and Temporal Contextualized Knowledge Embeddings for Temporal Knowledge Graph Completion</i> Yifu Gao, Yongquan He, Zhigang Kan, Yi Han, Linbo Qiao and Dongsheng Li	417
<i>A Systematic Study and Comprehensive Evaluation of ChatGPT on Benchmark Datasets</i> Md Tahmid Rahman Laskar, M Saiful Bari, Mizanur Rahman, Md Amran Hossen Bhuiyan, Shafiq Joty and Jimmy Xiangji Huang	431
<i>Generating Deep Questions with Commonsense Reasoning Ability from the Text by Disentangled Adversarial Inference</i> Jianxing Yu, Shiqi Wang, Libin Zheng, Qinliang Su, Wei Liu, Baoquan Zhao and Jian Yin ..	470
<i>TADA: Efficient Task-Agnostic Domain Adaptation for Transformers</i> Chia-Chien Hung, Lukas Lange and Jannik Strötgen	487
<i>Robust Natural Language Understanding with Residual Attention Debiasing</i> Fei Wang, James Y. Huang, Tianyi Yan, Wenxuan Zhou and Muhao Chen	504
<i>MoNET: Tackle State Momentum via Noise-Enhanced Training for Dialogue State Tracking</i> Haoning Zhang, Junwei Bao, Haipeng Sun, Youzheng Wu, Wenye Li, Shuguang Cui and Xiaodong He	520

<i>PAL: Persona-Augmented Emotional Support Conversation Generation</i>	
Jiale Cheng, Sahand Sabour, Hao Sun, Zhuang Chen and Minlie Huang	535
<i>Farewell to Aimless Large-scale Pretraining: Influential Subset Selection for Language Model</i>	
Xiao Wang, Weikang Zhou, Qi Zhang, Jie Zhou, Song Yang Gao, Junzhe Wang, Menghan Zhang, Xiang Gao, Yun Wen Chen and Tao Gui	555
<i>Exclusive Supermask Subnetwork Training for Continual Learning</i>	
Prateek Yadav and Mohit Bansal	569
<i>Transferring General Multimodal Pretrained Models to Text Recognition</i>	
Junyang Lin, Xuancheng Ren, Yichang Zhang, Gao Liu, Peng Wang, An Yang and Chang Zhou	588
<i>A Formal Perspective on Byte-Pair Encoding</i>	
Vilém Zouhar, Clara Meister, Juan Luis Gastaldi, Li Du, Tim Vieira, Mrinmaya Sachan and Ryan Cotterell	598
<i>Automatic Named Entity Obfuscation in Speech</i>	
Judita Preiss	615
<i>Recursion of Thought: A Divide-and-Conquer Approach to Multi-Context Reasoning with Language Models</i>	
Soochan Lee and Gunhee Kim	623
<i>UniS-MMC: Multimodal Classification via Unimodality-supervised Multimodal Contrastive Learning</i>	
Heqing Zou, Meng Shen, Chen Chen, Yuchen Hu, Deepu Rajan and Eng Siong Chng	659
<i>Robustness-Aware Word Embedding Improves Certified Robustness to Adversarial Word Substitutions</i>	
Yibin Wang, Yichen Yang, Di He and Kun He	673
<i>Exploring the Compositional Generalization in Context Dependent Text-to-SQL Parsing</i>	
Aiwei Liu, Wei Liu, Xuming Hu, Shuang Li, Fukun Ma, Yawen Yang and Lijie Wen	688
<i>Towards Generative Event Factuality Prediction</i>	
John Murzaku, Tyler G Osborne, Amittai F Aviram and Owen Rambow	701
<i>Can Language Models Be Specific? How?</i>	
Jie Huang, Kevin Chen-Chuan Chang, Jinjun Xiong and Wen-mei Hwu	716
<i>The Web Can Be Your Oyster for Improving Language Models</i>	
Junyi Li, Tianyi Tang, Wayne Xin Zhao, Jingyuan Wang, Jian-Yun Nie and Ji-Rong Wen	728
<i>Enhancing Few-shot Cross-lingual Transfer with Target Language Peculiar Examples</i>	
Hwichan Kim and Mamoru Komachi	747
<i>Overcoming Catastrophic Forgetting in Massively Multilingual Continual Learning</i>	
Genta Indra Winata, Lingjue Xie, Karthik Radhakrishnan, Shijie Wu, Xisen Jin, Pengxiang Cheng, Mayank Kulkarni and Daniel Preotiuc-Pietro	768
<i>UniFine: A Unified and Fine-grained Approach for Zero-shot Vision-Language Understanding</i>	
Rui Sun, Zhecan Wang, Haoxuan You, Noel Codella, Kai-Wei Chang and Shih-Fu Chang ...	778
<i>Aligning Instruction Tasks Unlocks Large Language Models as Zero-Shot Relation Extractors</i>	
Kai Zhang, Bernal Jimenez Gutierrez and Yu Su	794
<i>TADA : Task Agnostic Dialect Adapters for English</i>	
William Held, Caleb Ziems and Diyi Yang	813

<i>Generative Zero-Shot Prompt Learning for Cross-Domain Slot Filling with Inverse Prompting</i> Xuefeng Li, Liwen Wang, Guanting Dong, Keqing He, Jinzheng Zhao, Hao Lei, Jiachi Liu and Weiran Xu	825
<i>Re-appraising the Schema Linking for Text-to-SQL</i> Yujian Gan, Xinyun Chen and Matthew Purver	835
<i>Echoes from Alexandria: A Large Resource for Multilingual Book Summarization</i> Alessandro Scirè, Simone Conia, Simone Ciciliano and Roberto Navigli	853
<i>When Gradient Descent Meets Derivative-Free Optimization: A Match Made in Black-Box Scenario</i> Chengcheng Han, Liqing Cui, Renyu Zhu, Jianing Wang, Nuo Chen, Qiushi Sun, Xiang Li and Ming Gao	868
<i>Align-then-Enhance: Multilingual Entailment Graph Enhancement with Soft Predicate Alignment</i> Yuting Wu, Yutong Hu, Yansong Feng, Tianyi Li, Mark Steedman and Dongyan Zhao	881
<i>Few-shot Classification with Hypersphere Modeling of Prototypes</i> Ning Ding, Yulin Chen, Ganqu CUI, Xiaobin Wang, Haitao Zheng, Zhiyuan Liu and Pengjun Xie	895
<i>Structured Mean-Field Variational Inference for Higher-Order Span-Based Semantic Role</i> Wei Liu, Songlin Yang and Kewei Tu	918
<i>AQE: Argument Quadruplet Extraction via a Quad-Tagging Augmented Generative Approach</i> Jia Guo, Liying Cheng, Wenxuan Zhang, Stanley Kok, Xin Li and Lidong Bing	932
<i>The Dangers of trusting Stochastic Parrots: Faithfulness and Trust in Open-domain Conversational Question Answering</i> Sabrina Chiesurin, Dimitris Dimakopoulos, Marco Antonio Sobrevilla Cabezudo, Arash Eshghi, Ioannis Papaioannou, Verena Rieser and Ioannis Konstas	947
<i>Discrete Prompt Optimization via Constrained Generation for Zero-shot Re-ranker</i> Sukmin Cho, Soyeong Jeong, Jeong yeon Seo and Jong Park	960
<i>Triggering Multi-Hop Reasoning for Question Answering in Language Models using Soft Prompts and Random Walks</i> Kanishka Misra, Cicero Nogueira dos Santos and Siamak Shakeri	972
<i>Multimedia Generative Script Learning for Task Planning</i> Qingyun Wang, Manling Li, Hou Pong Chan, Lifu Huang, Julia Hockenmaier, Girish Chowdhary and Heng Ji	986
<i>Label Agnostic Pre-training for Zero-shot Text Classification</i> Christopher Clarke, Yuzhao Heng, Yiping Kang, Krisztian Flautner, Lingjia Tang and Jason Mars	1009
<i>Click: Controllable Text Generation with Sequence Likelihood Contrastive Learning</i> Chujie Zheng, Pei Ke, Zheng Zhang and Minlie Huang	1022
<i>Improving Embedding-based Unsupervised Keyphrase Extraction by Incorporating Structural Information</i> Mingyang Song, Huafeng Liu, Yi Feng and Liping Jing	1041
<i>Towards Reasoning in Large Language Models: A Survey</i> Jie Huang and Kevin Chen-Chuan Chang	1049

<i>Transitioning from benchmarks to a real-world case of information-seeking in Scientific Publications</i> Chyrine Tahri, Aurore Bochnakian, Patrick Haouat and Xavier Tannier	1066
<i>CLIPText: A New Paradigm for Zero-shot Text Classification</i> Libo Qin, Weiyun Wang, Qiguang Chen and Wanxiang Che	1077
<i>Rethinking Dictionaries and Glyphs for Chinese Language Pre-training</i> Yuxuan Wang, Jack Wang, Dongyan Zhao and Zilong Zheng	1089
<i>One Embedder, Any Task: Instruction-Finetuned Text Embeddings</i> Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer and Tao Yu	1102
<i>Towards Speech Dialogue Translation Mediating Speakers of Different Languages</i> Shuichiro Shimizu, Chenhui Chu, Sheng Li and Sadao Kurohashi	1122
<i>Adaptation Approaches for Nearest Neighbor Language Models</i> Rishabh Bhardwaj, George Polovets and Monica Sunkara	1135
<i>Language Models for German Text Simplification: Overcoming Parallel Data Scarcity through Style-specific Pre-training</i> Miriam Anschutz, Joshua Oehms, Thomas Wimmer, Bartłomiej Jezierski and Georg Groh .	1147
<i>Client-Customized Adaptation for Parameter-Efficient Federated Learning</i> Yeachen Kim, Junho Kim, Wing-Lam Mok, Jun-Hyung Park and SangKeun Lee	1159
<i>FolkScope: Intention Knowledge Graph Construction for E-commerce Commonsense Discovery</i> Changlong Yu, Weiqi Wang, Xin Liu, Jiabin Bai, Yangqiu Song, Zheng Li, Yifan Gao, Tianyu Cao and Bing Yin	1173
<i>I am PsyAM: Modeling Happiness with Cognitive Appraisal Dimensions</i> Xuan Liu and Kokil Jaidka	1192
<i>Value type: the bridge to a better DST model</i> Gao Qixiang, Mingyang Sun, Yutao Mou, Chen Zeng and Weiran Xu	1211
<i>Hypothetical Training for Robust Machine Reading Comprehension of Tabular Context</i> Moxin Li, Wenjie Wang, Fuli Feng, Hanwang Zhang, Qifan Wang and Tat-Seng Chua	1220
<i>BanglaBook: A Large-scale Bangla Dataset for Sentiment Analysis from Book Reviews</i> Mohsinul Kabir, Obayed Bin Mahfuz, Syed Rifat Raiyan, Hasan Mahmud and Md Kamrul Hasan	1237
<i>Risks and NLP Design: A Case Study on Procedural Document QA</i> Nikita Haduong, Alice Gao and Noah A. Smith	1248
<i>The Diminishing Returns of Masked Language Models to Science</i> Zhi Hong, Aswathy Ajith, James Gregory Pauloski, Eamon Duede, kyle chard and Ian T Foster	1270
<i>Causal Matching with Text Embeddings: A Case Study in Estimating the Causal Effects of Peer Review Policies</i> Raymond Zhang, Neha Nayak Kennard, Daniel S Smith, Daniel A McFarland, Andrew McCallum and Katherine A Keith	1284
<i>Learning to Generalize for Cross-domain QA</i> Yingjie Niu, Linyi Yang, Ruihai Dong and Yue Zhang	1298

<i>Enhanced Chart Understanding via Visual Language Pre-training on Plot Table Pairs</i>	
Mingyang Zhou, Yi Fung, Long Chen, Christopher Thomas, Heng Ji and Shih-Fu Chang . . .	1314
<i>Importance of Synthesizing High-quality Data for Text-to-SQL Parsing</i>	
Yiqun Hu, Yiyun Zhao, Jiarong Jiang, Wuwei Lan, Henghui Zhu, Anuj Chauhan, Alexander Hanbo Li, Lin Pan, Jun Wang, Chung-Wei Hang, Sheng Zhang, Jiang Guo, Mingwen Dong, Joseph Lilien, Patrick Ng, Zhiguo Wang, Vittorio Castelli and Bing Xiang	1327
<i>Exploring Schema Generalizability of Text-to-SQL</i>	
Jieyu Li, Lu Chen, Ruisheng Cao, Su Zhu, Hongshen Xu, Zhi Chen, Hanchong Zhang and Kai Yu	1344
<i>Enhancing Cross-lingual Natural Language Inference by Soft Prompting with Multilingual Verbalizer</i>	
Shuang Li, Xuming Hu, Aiwei Liu, Yawen Yang, Fukun Ma, Philip S. Yu and Lijie Wen . . .	1361
<i>A Confidence-based Partial Label Learning Model for Crowd-Annotated Named Entity Recognition</i>	
Limao Xiong, Jie Zhou, Qunxi Zhu, Xiao Wang, Yuanbin Wu, Qi Zhang, Tao Gui, Xuanjing Huang, Jin Ma and Ying Shan	1375
<i>Towards Zero-Shot Persona Dialogue Generation with In-Context Learning</i>	
Xinchao Xu, Zeyang Lei, Wenquan Wu, Zheng-Yu Niu, Hua Wu and Haifeng Wang	1387
<i>Grammar-based Decoding for Improved Compositional Generalization in Semantic Parsing</i>	
Jing Zheng, Jyh-Herng Chow, Zhongnan Shen and Peng Xu	1399
<i>Exploiting Rich Textual User-Product Context for Improving Personalized Sentiment Analysis</i>	
Chenyang Lyu, Linyi Yang, Yue Zhang, Yvette Graham and Jennifer Foster	1419
<i>Efficient Out-of-Domain Detection for Sequence to Sequence Models</i>	
Artem Vazhentsev, Akim Tsvigun, Roman Konstantinovich Vashurin, Sergey Petrakov, Daniil Vasilev, Maxim Panov, Alexander Panchenko and Artem Shelmanov	1430
<i>Emotion Cause Extraction on Social Media without Human Annotation</i>	
Debin Xiao, Rui Xia and Jianfei Yu	1455
<i>Pseudo Outlier Exposure for Out-of-Distribution Detection using Pretrained Transformers</i>	
Jaeyoung Kim, Kyuheon Jung, Dongbin Na, Sion Jang, Eunbin Park and Sungchul Choi . . .	1469
<i>Adversarial Multi-task Learning for End-to-end Metaphor Detection</i>	
Shenglong Zhang and Ying Liu	1483
<i>SERENGETI: Massively Multilingual Language Models for Africa</i>	
Ife Adebara, AbdelRahim Elmadany, Muhammad Abdul-Mageed and Alcides Alcoba Inciarte	1498
<i>Prompt- and Trait Relation-aware Cross-prompt Essay Trait Scoring</i>	
Heejin Do, Yunsu Kim and Gary Geunbae Lee	1538
<i>AugESC: Dialogue Augmentation with Large Language Models for Emotional Support Conversation</i>	
Chujie Zheng, Sahand Sabour, Jiaxin Wen, Zheng Zhang and Minlie Huang	1552
<i>2 * n is better than n²: Decomposing Event Coreference Resolution into Two Tractable Problems</i>	
Shafiuddin Rehan Ahmed, Abhijnan Nath, James H. Martin and Nikhil Krishnaswamy	1569
<i>SCCS: Semantics-Consistent Cross-domain Summarization via Optimal Transport Alignment</i>	
Jielin Qiu, Jiacheng Zhu, Mengdi Xu, Franck Dernoncourt, Trung Bui, Zhaowen Wang, Bo Li, Ding Zhao and Hailin Jin	1584

<i>General-to-Specific Transfer Labeling for Domain Adaptable Keyphrase Generation</i>	
Rui Meng, Tong Wang, Xingdi Yuan, Yingbo Zhou and Daqing He	1602
<i>E-NER: Evidential Deep Learning for Trustworthy Named Entity Recognition</i>	
Zhen Zhang, Mengting Hu, Shiwang Zhao, Minlie Huang, Haotian Wang, Lemao Liu, Zhirui Zhang, Zhe Liu and Bingzhe Wu	1619
<i>LMCap: Few-shot Multilingual Image Captioning by Retrieval Augmented Language Model Prompting</i>	
Rita Parada Ramos, Bruno Martins and Desmond Elliott	1635
<i>Boosting Text Augmentation via Hybrid Instance Filtering Framework</i>	
Heng Yang and Ke Li	1652
<i>Gradient-Boosted Decision Tree for Listwise Context Model in Multimodal Review Helpfulness Prediction</i>	
Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy T. Nguyen, Zhen Hai, Lidong Bing and Anh Tuan Luu	1670
<i>Extract and Attend: Improving Entity Translation in Neural Machine Translation</i>	
Zixin Zeng, Rui Wang, Yichong Leng, Junliang Guo, Shufang Xie, Xu Tan, Tao Qin and Tie-Yan Liu	1697
<i>Real-World Compositional Generalization with Disentangled Sequence-to-Sequence Learning</i>	
Hao Zheng and Mirella Lapata	1711
<i>Cross-lingual AMR Aligner: Paying Attention to Cross-Attention</i>	
Abelardo Carlos Martínez Lorenzo, Pere Lluís Hugué Cabot and Roberto Navigli	1726
<i>Zero-Shot Text Classification via Self-Supervised Tuning</i>	
Chaoqun Liu, Wenxuan Zhang, Guizhen Chen, Xiaobao Wu, Anh Tuan Luu, Chip Hong Chang and Lidong Bing	1743
<i>Logical Transformers: Infusing Logical Structures into Pre-Trained Language Models</i>	
Borui Wang, Qiuyuan Huang, Budhaditya Deb, Aaron L Halfaker, Liqun Shao, Daniel McDuff, Ahmed Hassan Awadallah, Dragomir Radev and Jianfeng Gao	1762
<i>Large Language Models with Controllable Working Memory</i>	
Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu and Sanjiv Kumar	1774
<i>A Unified Evaluation Framework for Novelty Detection and Accommodation in NLP with an Instantiation in Authorship Attribution</i>	
Neeraj Varshney, Himanshu Gupta, Eric Robertson, Bing Liu and Chitta Baral	1794
<i>CDA: A Contrastive Data Augmentation Method for Alzheimer’s Disease Detection</i>	
Junwen Duan, Fangyuan Wei, Jin Liu, Hongdong Li, Tianming Liu and Jianxin Wang	1819
<i>Disentangling Aspect and Stance via a Siamese Autoencoder for Aspect Clustering of Vaccination Opinions</i>	
Lixing Zhu, Runcong Zhao, Gabriele Pergola and Yulan He	1827
<i>Temporal Relation Classification using Boolean Question Answering</i>	
Omer Cohen and Kfir Bar	1843
<i>Are Synonym Substitution Attacks Really Synonym Substitution Attacks?</i>	
Cheng-Han Chiang and Hung-yi Lee	1853

<i>DivHSK: Diverse Headline Generation using Self-Attention based Keyword Selection</i>	
Venkatesh E, Kaushal Kumar Maurya, Deepak Kumar and Maunendra Sankar Desarkar . . .	1879
<i>Similarity-Based Content Scoring - A more Classroom-Suitable Alternative to Instance-Based Scoring?</i>	
Marie Bexte, Andrea Horbach and Torsten Zesch	1892
<i>Pragmatic Inference with a CLIP Listener for Contrastive Captioning</i>	
Jiefu Ou, Benno Krojer and Daniel Fried	1904
<i>A Statistical Exploration of Text Partition Into Constituents: The Case of the Priestly Source in the Books of Genesis and Exodus</i>	
Gideon Gidi Yoffe, Axel Bühler, Nachum Dershowitz, Thomas C Romer, Eli Piasetzky, Israel Finkelstein and Barak Sober	1918
<i>A Language-First Approach for Procedure Planning</i>	
Jiateng Liu, Sha Li, Zhenhailong Wang, Manling Li and Heng Ji	1941
<i>An Empirical Analysis of Leveraging Knowledge for Low-Resource Task-Oriented Semantic Parsing</i>	
Mayank Kulkarni, Aoxiao Zhong, Nicolas Guenon des mesnards, Sahar Movaghati, Mukund Sridhar, He Xie and Jianhua Lu	1955
<i>TempLM: Distilling Language Models into Template-Based Generators</i>	
Tianyi Zhang, Mina Lee, Xiang Lisa Li, Ende Shen and Tatsunori Hashimoto	1970
<i>Incorporating Graph Information in Transformer-based AMR Parsing</i>	
Pavlo Vasylenko, Pere Lluís Huguet Cabot, Abelardo Carlos Martínez Lorenzo and Roberto Navigli	1995
<i>Rethinking the Word-level Quality Estimation for Machine Translation from Human Judgement</i>	
Zhen Yang, Fandong Meng, Yuanmeng Yan and Jie Zhou	2012
<i>PV2TEA: Patching Visual Modality to Textual-Established Information Extraction</i>	
Hejie Cui, Rongmei Lin, Nasser Zalmout, Chenwei Zhang, Jingbo Shang, Carl Yang and Xian Li	2026
<i>Structural Contrastive Pretraining for Cross-Lingual Comprehension</i>	
Nuo Chen, Linjun Shou, Tengtao Song, Ming Gong, Jian Pei, Jianhui Chang, Daxin Jiang and Jia Li	2042
<i>Reducing Sensitivity on Speaker Names for Text Generation from Dialogues</i>	
Qi Jia, Haifeng Tang and Kenny Q. Zhu	2058
<i>Topic and Style-aware Transformer for Multimodal Emotion Recognition</i>	
Shuwen Qiu, Nitesh Sekhar and Prateek Singhal	2074
<i>Exploiting Abstract Meaning Representation for Open-Domain Question Answering</i>	
Cunxiang Wang, Zhikun Xu, Qipeng Guo, Xiangkun Hu, Xuefeng Bai, Zheng Zhang and Yue Zhang	2083
<i>Nonparametric Masked Language Modeling</i>	
Sewon Min, Weijia Shi, Mike Lewis, Xilun Chen, Wen-tau Yih, Hannaneh Hajishirzi and Luke Zettlemoyer	2097
<i>Pay More Attention to Relation Exploration for Knowledge Base Question Answering</i>	
Yong Cao, Xianzhi Li, huiwen liu, Wen Dai, shuai chen, Bin Wang, Min Chen and Daniel Hershcovich	2119

<i>Speaking Multiple Languages Affects the Moral Bias of Language Models</i> Katharina Haemmerl, Bjoern Deiseroth, Patrick Schramowski, Jindřich Libovický, Constantin A Rothkopf, Alexander Fraser and Kristian Kersting	2137
<i>Retrieving Relevant Context to Align Representations for Cross-lingual Event Detection</i> Chien Van Nguyen, Linh Ngo and Thien Huu Nguyen	2157
<i>NormNet: Normalize Noun Phrases for More Robust NLP</i> Minlong Peng and Mingming Sun	2171
<i>Cross Encoding as Augmentation: Towards Effective Educational Text Classification</i> Hyun Seung Lee, Seungtaek Choi, Yunsung Lee, Hyeongdon Moon, Shinhyeok Oh, Myeongho Jeong, Hyojun Go and Christian Wallraven	2184
<i>Adversarial Robustness of Prompt-based Few-Shot Learning for Natural Language Understanding</i> Venkata Prabhakara Sarath Nookala, Gaurav Verma, Subhabrata Mukherjee and Srijan Kumar	2196
<i>This prompt is measuring <mask>: evaluating bias evaluation in language models</i> Seraphina Goldfarb-Tarrant, Eddie L. Ungless, Esma Balkir and Su Lin Blodgett	2209
<i>Towards Open Environment Intent Prediction</i> Yunhua Zhou, Jiawei Hong and Xipeng Qiu	2226
<i>Teamwork Is Not Always Good: An Empirical Study of Classifier Drift in Class-incremental Information Extraction</i> Minqian Liu and Lifu Huang	2241
<i>C-XNLI: Croatian Extension of XNLI Dataset</i> Leo Obadić, Andrej Jertec, Marko Rajnović and Branimir Dropuljić	2258
<i>AVATAR: A Parallel Corpus for Java-Python Program Translation</i> Wasi Uddin Ahmad, Md Golam Rahman Tushar, Saikat Chakraborty and Kai-Wei Chang ..	2268
<i>On Dataset Transferability in Active Learning for Transformers</i> Fran Jelenić, Josip Jukić, Nina Drobac and Jan Snajder	2282
<i>Structured Persuasive Writing Support in Legal Education: A Model and Tool for German Legal Case Solutions</i> Florian Weber, Thiemo Wambsganss, Seyed Parsa Neshaei and Matthias Soellner	2296
<i>Characterizing the Impacts of Instances on Robustness</i> Rui Zheng, Zhiheng Xi, Qin Liu, Wenbin Lai, Tao Gui, Qi Zhang, Xuanjing Huang, Jin Ma, Ying Shan and Weifeng Ge	2314
<i>Generate then Select: Open-ended Visual Question Answering Guided by World Knowledge</i> Xingyu Fu, Sheng Zhang, Gukyeong Kwon, Pramuditha Perera, Henghui Zhu, Yuhao Zhang, Alexander Hanbo Li, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Patrick Ng, Dan Roth and Bing Xiang	2333
<i>Hence, Socrates is mortal: A Benchmark for Natural Language Syllogistic Reasoning</i> Yongkang Wu, meng han, Yutao Zhu, Lei Li, Xinyu Zhang, ruofei lai, Xiaoguang Li, Yuanhang Ren, Zhicheng Dou and Zhao Cao	2347
<i>Categorical grammar induction from raw data</i> Christian Clark and William Schuler	2368

<i>Attribute Controlled Dialogue Prompting</i>	
Runcheng Liu, Ahmad Rashid, Ivan Kobzyev, Mehdi Rezagholizadeh and Pascal Poupart . . .	2380
<i>Open-World Factually Consistent Question Generation</i>	
Himanshu Maheshwari, Sumit Shekhar, Apoorv Saxena and Niyati Chhaya	2390
<i>Contrastive Learning of Sociopragmatic Meaning in Social Media</i>	
Chiyu Zhang, Muhammad Abdul-Mageed and Ganesh Jawahar	2405
<i>Noisy Positive-Unlabeled Learning with Self-Training for Speculative Knowledge Graph Reasoning</i>	
Ruijie Wang, Baoyu Li, Yichen Lu, Dachun Sun, Jinning Li, Yuchen Yan, Shengzhong Liu, Hanghang Tong and Tarek F. Abdelzaher	2440
<i>ACROSS: An Alignment-based Framework for Low-Resource Many-to-One Cross-Lingual Summarization</i>	
Peiyao Li, Zhengkun Zhang, Jun Wang, Liang Li, Adam Jatowt and Zhenglu Yang	2458
<i>RFiD: Towards Rational Fusion-in-Decoder for Open-Domain Question Answering</i>	
Cunxiang Wang, Haofei Yu and Yue Zhang	2473
<i>Unsupervised Keyphrase Extraction by Learning Neural Keyphrase Set Function</i>	
Mingyang Song, Haiyun Jiang, Lemaou Liu, Shuming Shi and Liping Jing	2482
<i>Diffusion Theory as a Scalpel: Detecting and Purifying Poisonous Dimensions in Pre-trained Language Models Caused by Backdoor or Bias</i>	
Zhiyuan Zhang, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou and Xu Sun	2495
<i>Retrieving Multimodal Prompts for Generative Visual Question Answering</i>	
Timothy Ossowski and Junjie Hu	2518
<i>InfoSync: Information Synchronization across Multilingual Semi-structured Tables</i>	
Siddharth Hemant Khincha, Chelsi Jain, Vivek Gupta, Tushar Kataria and Shuo Zhang	2536
<i>T2IAT: Measuring Valence and Stereotypical Biases in Text-to-Image Generation</i>	
Jialu Wang, Xinyue Gabby Liu, Zonglin Di, Yang Liu and Xin Eric Wang	2560
<i>An Investigation of Evaluation Methods in Automatic Medical Note Generation</i>	
Asma Ben Abacha, Wen-wai Yim, George Michalopoulos and Thomas Lin	2575
<i>Rethinking Translation Memory Augmented Neural Machine Translation</i>	
Hongkun Hao, Guoping Huang, Lemaou Liu, Zhirui Zhang, Shuming Shi and Rui Wang	2589
<i>Controlling Styles in Neural Machine Translation with Activation Prompt</i>	
Yifan Wang, Zewei Sun, Shanbo Cheng, Weiguo Zheng and Mingxuan Wang	2606
<i>Focusing, Bridging and Prompting for Few-shot Nested Named Entity Recognition</i>	
Yuanyuan Xu, Zeng Yang, Linhai Zhang, Deyu Zhou, Tiandeng Wu and Rong Zhou	2621
<i>Together We Make Sense—Learning Meta-Sense Embeddings</i>	
Haochen Luo, Yi Zhou and Danushka Bollegala	2638
<i>Multimodal Prompt Learning for Product Title Generation with Extremely Limited Labels</i>	
Bang Yang, Fenglin Liu, Zheng Li, Qingyu Yin, Chenyu You, Bing Yin and Yuexian Zou	2652
<i>Large Language Models are Built-in Autoregressive Search Engines</i>	
Noah Ziemis, Wenhao Yu, Zhihan Zhang and Meng Jiang	2666

<i>Beyond Triplet: Leveraging the Most Data for Multimodal Machine Translation</i>	
Yaoming ZHU, Zewei Sun, Shanbo Cheng, Luyang Huang, Liwei Wu and Mingxuan Wang	2679
<i>From chocolate bunny to chocolate crocodile: Do Language Models Understand Noun Compounds?</i>	
Albert Jordan Coil and Vered Shwartz	2698
<i>Measuring Intersectional Biases in Historical Documents</i>	
Nadav Borenstein, Karolina Stanczak, Thea Rolskov, Natacha Klein Käfer, Natália da Silva Perez and Isabelle Augenstein	2711
<i>Incomplete Utterance Rewriting by A Two-Phase Locate-and-Fill Regime</i>	
Zitong Li, Jiawei Li, Haifeng Tang, Kenny Q. Zhu and Ruolan Yang	2731
<i>Exploring Variation of Results from Different Experimental Conditions</i>	
Maja Popović, Mohammad Arvan, Natalie Parde and Anya Belz	2746
<i>Playing the Part of the Sharp Bully: Generating Adversarial Examples for Implicit Hate Speech Detection</i>	
Nicolas Benjamin Ocampo, Elena Cabrio and Serena Villata	2758
<i>X-RiSAWOZ: High-Quality End-to-End Multilingual Dialogue Datasets and Few-shot Agents</i>	
Mehrad Moradshahi, Tianhao Shen, Kalika Bali, Monojit Choudhury, Gael de Chalendar, Anmol Goel, Sungkyun Kim, Prashant Kodali, Ponnurangam Kumaraguru, Nasredine Semmar, Sina Semnani, Jiwon Seo, Vivek Seshadri, Manish Shrivastava, Michael Sun, Aditya Yadavalli, Chaobin You, Deyi Xiong and Monica S Lam	2773
<i>Subword Segmental Machine Translation: Unifying Segmentation and Target Sentence Generation</i>	
Francois Meyer and Jan Buys	2795
<i>Measuring and Mitigating Local Instability in Deep Neural Networks</i>	
Arghya Datta, Subhrangshu Nandi, Jingcheng Xu, Greg Ver Steeg, He Xie, Anoop Kumar and Aram Galstyan	2810
<i>What Knowledge Is Needed? Towards Explainable Memory for kNN-MT Domain Adaptation</i>	
Wenhao Zhu, Shujian Huang, Yunzhe Lv, Xin Zheng and Jiajun CHEN	2824
<i>Measuring Your ASTE Models in The Wild: A Diversified Multi-domain Dataset For Aspect Sentiment Triplet Extraction</i>	
Ting Xu, Huiyun Yang, Zhen Wu, Jiase Chen, Fei Zhao and Xinyu Dai	2837
<i>Grounding the Lexical Substitution Task in Entailment</i>	
Talgat Omarov and Grzegorz Kondrak	2854
<i>Operator Selection and Ordering in a Pipeline Approach to Efficiency Optimizations for Transformers</i>	
Ji Xin, Raphael Tang, Zhiying Jiang, Yaoliang Yu and Jimmy Lin	2870
<i>AraMUS: Pushing the Limits of Data and Model Scale for Arabic Natural Language Processing</i>	
Asaad Alghamdi, Xinyu Duan, Wei Jiang, Zhenhai Wang, Yimeng Wu, Qingrong Xia, Zhefeng Wang, Yi ZHENG, Mehdi Rezagholizadeh, baoxing Huai, Peilun Cheng and Abbas Ghaddar	2883
<i>Leveraging Explicit Procedural Instructions for Data-Efficient Action Prediction</i>	
Julia Isabel White, Arushi Raghuvanshi and Yada Pruksachatkun	2895
<i>Quantifying Train-Evaluation Overlap with Nearest Neighbors</i>	
Gauri Kambhatla, Thuy H Nguyen and Eunsol Choi	2905
<i>Unsupervised Mapping of Arguments of Deverbal Nouns to Their Corresponding Verbal Labels</i>	
Aviv Weinstein and Yoav Goldberg	2921

<i>The Decades Progress on Code-Switching Research in NLP: A Systematic Survey on Trends and Challenges</i>	
Genta Indra Winata, Alham Fikri Aji, Zheng Xin Yong and Thamar Solorio	2936
<i>Learning to Predict Persona Information for Dialogue Personalization without Explicit Persona Description</i>	
Wangchunshu Zhou, Qifei LI and Chenle Li	2979
<i>Automated Refugee Case Analysis: A NLP Pipeline for Supporting Legal Practitioners</i>	
Claire Barale, Michael Rovatsos and Nehal Bhuta	2992
<i>Recurrent Attention Networks for Long-text Modeling</i>	
Xianming Li, Zongxi Li, Xiaotian Luo, Haoran Xie, Xing Lee, Yingbin Zhao, Fu Lee Wang and Qing Li	3006
<i>Exploring the Relationship between Alignment and Cross-lingual Transfer in Multilingual Transformers</i>	
Felix Gaschi, Patricio Cerda, Parisa RASTIN and Yannick Toussaint	3020
<i>Aerial Vision-and-Dialog Navigation</i>	
Yue Fan, Winson Chen, Tongzhou Jiang, Chun Zhou, Yi Zhang and Xin Eric Wang	3043
<i>Improved Logical Reasoning of Language Models via Differentiable Symbolic Programming</i>	
Hanlin Zhang, Jiani Huang, Ziyang Li, MAYUR NAIK and Eric Xing	3062
<i>B2T Connection: Serving Stability and Performance in Deep Transformers</i>	
Sho Takase, Shun Kiyono, Sosuke Kobayashi and Jun Suzuki	3078
<i>Boosting Zero-shot Cross-lingual Retrieval by Training on Artificially Code-Switched Data</i>	
Robert Litschko, Ekaterina Artemova and Barbara Plank	3096
<i>Domain-specific Attention with Distributional Signatures for Multi-Domain End-to-end Task-Oriented Dialogue</i>	
Xing Ma, Peng Zhang and Feifei Zhao	3109
<i>CKDST: Comprehensively and Effectively Distill Knowledge from Machine Translation to End-to-End Speech Translation</i>	
Yikun Lei, Zhengshan Xue, Xiaohu Zhao, Haoran Sun, Shaolin Zhu, xiaodong lin and Deyi Xiong	3123
<i>Follow the leader(board) with confidence: Estimating p-values from a single test set with item and response variance</i>	
Shira Wein, Christopher Homan, Lora Aroyo and Chris Welty	3138
<i>Parallel Data Helps Neural Entity Coreference Resolution</i>	
Gongbo Tang and Christian Hardmeier	3162
<i>Towards Open-Domain Twitter User Profile Inference</i>	
Haoyang Wen, Zhenxin Xiao, Eduard H Hovy and Alexander Hauptmann	3172
<i>Eliciting Affective Events from Language Models by Multiple View Co-prompting</i>	
Yuan Zhuang and Ellen Riloff	3189
<i>ZeroAE: Pre-trained Language Model based Autoencoder for Transductive Zero-shot Text Classification</i>	
Kaihao Guo, Hang Yu, Cong Liao, Jianguo Li and Haipeng Zhang	3202

<i>PRAM: An End-to-end Prototype-based Representation Alignment Model for Zero-resource Cross-lingual Named Entity Recognition</i>	
Yucheng Huang, wenqiang liu, Xianli Zhang, JUN LANG, Tieliang Gong and Chen Li	3220
<i>It Takes Two to Tango: Navigating Conceptualizations of NLP Tasks and Measurements of Performance</i>	
Arjun Subramonian, Xingdi Yuan, Hal Daumé III and Su Lin Blodgett	3234
<i>Task-adaptive Label Dependency Transfer for Few-shot Named Entity Recognition</i>	
Shan Zhang, Bin Cao, Tianming Zhang, Yuqi Liu and Jing Fan	3280
<i>WYWEB: A NLP Evaluation Benchmark For Classical Chinese</i>	
Bo Zhou, Qianglong Chen, Xiaomi Zhong, Xiaomi Zhong and Yin Zhang	3294
<i>A Fused Gromov-Wasserstein Framework for Unsupervised Knowledge Graph Entity Alignment</i>	
Jianheng Tang, Kangfei Zhao and Jia Li	3320
<i>Two Examples are Better than One: Context Regularization for Gradient-based Prompt Tuning</i>	
Hyeonmin Ha, SOYOUNG JUNG, Jinsol Park, Minjoon Seo, Seung-won Hwang and Byung-Gon Chun	3335
<i>An Investigation of Noise in Morphological Inflection</i>	
Adam Wiemerslage, Changbing Yang, Garrett Nicolai, Miikka Silfverberg and Katharina Kann	3351
<i>Graph Reasoning for Question Answering with Triplet Retrieval</i>	
Shiyang Li, Yifan Gao, Haoming Jiang, Qingyu Yin, Zheng Li, Xifeng Yan, Chao Zhang and Bing Yin	3366
<i>End-to-End Argument Mining over Varying Rhetorical Structures</i>	
Elena Chistova	3376
<i>Unsupervised Task Graph Generation from Instructional Video Transcripts</i>	
Lajanugen Logeswaran, Sungryull Sohn, Yunseok Jang, Moontae Lee and Honglak Lee	3392
<i>Exploiting Hierarchically Structured Categories in Fine-grained Chinese Named Entity Recognition</i>	
Jiuding Yang, Jinwen Luo, Weidong Guo, Di Niu and Yu Xu	3407
<i>Adversarial Textual Robustness on Visual Dialog</i>	
Lu Yu and Verena Rieser	3422
<i>Language Model Analysis for Ontology Subsumption Inference</i>	
Yuan He, Jiaoyan Chen, Ernesto Jimenez-Ruiz, Hang Dong and Ian Horrocks	3439
<i>Exploring Automatically Perturbed Natural Language Explanations in Relation Extraction</i>	
Wanyun Cui and Xingran Chen	3454
<i>Varta: A Large-Scale Headline-Generation Dataset for Indic Languages</i>	
Rahul Aralikkatte, Ziling Cheng, Sumanth Doddapaneni and Jackie Chi Kit Cheung	3468
<i>Better Zero-Shot Reasoning with Self-Adaptive Prompting</i>	
Kingchen Wan, Ruoxi SUN, Hanjun Dai, Sercan Arik and Tomas Pfister	3493
<i>Multimodal Recommendation Dialog with Subjective Preference: A New Challenge and Benchmark</i>	
Yuxing Long, Binyuan Hui, Caixia Yuan, Fei Huang, Yongbin Li and Xiaojie WANG	3515
<i>ANALOGICAL - A Novel Benchmark for Long Text Analogy Evaluation in Large Language Models</i>	
Thilini Wijesiriwardene, Ruwan Wickramarachchi, Bimal Gajera, Shreeyash Mukul Gowaikar, Chandan Gupta, Aman Chadha, Aishwarya Naresh Reganti, Amit Sheth and Amitava Das	3534

<i>Financial Numeric Extreme Labelling: A dataset and benchmarking</i>	
Soumya Sharma, Subhendu Khatuya, Manjunath Hegde, Afreen Shaikh, Koustuv Dasgupta, Pawan Goyal and Niloy Ganguly	3550
<i>Multilingual Summarization with Factual Consistency Evaluation</i>	
Roei Aharoni, Shashi Narayan, Joshua Maynez, Jonathan Herzig, Elizabeth Clark and Mirella Lapata	3562
<i>Enhancing Out-of-Vocabulary Estimation with Subword Attention</i>	
Raj Patel and Carlotta Domeniconi	3592
<i>Encoder and Decoder, Not One Less for Pre-trained Language Model Sponsored NMT</i>	
Sufeng Duan and Hai Zhao	3602
<i>TransGEC: Improving Grammatical Error Correction with Translationese</i>	
Tao Fang, Xuebo Liu, Derek F. Wong, Runzhe Zhan, Liang Ding, Lidia S. Chao, Dacheng Tao and Min Zhang	3614
<i>NewsDialogues: Towards Proactive News Grounded Conversation</i>	
Siheng Li, Yichun Yin, Cheng Yang, Wangjie Jiang, Yiwei Li, Zesen Cheng, Lifeng Shang, Xin Jiang, Qun Liu and Yujia Yang	3634
<i>Task-aware Retrieval with Instructions</i>	
Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi and Wen-tau Yih	3650
<i>Non-Repeatable Experiments and Non-Reproducible Results: The Reproducibility Crisis in Human Evaluation in NLP</i>	
Anya Belz, Craig Thomson, Ehud Reiter and Simon Mille	3676
<i>Define, Evaluate, and Improve Task-Oriented Cognitive Capabilities for Instruction Generation Models</i>	
Lingjun Zhao, Khanh Nguyen and Hal Daumé III	3688
<i>Robustness of Multi-Source MT to Transcription Errors</i>	
Dominik Macháček, Peter Polak, Ondřej Bojar and Raj Dabre	3707
<i>Not The End of Story: An Evaluation of ChatGPT-Driven Vulnerability Description Mappings</i>	
Xin Liu, Yuan Tan, Zhenghang Xiao, Jianwei Zhuge and Rui Zhou	3724
<i>Multi3NLU++: A Multilingual, Multi-Intent, Multi-Domain Dataset for Natural Language Understanding in Task-Oriented Dialogue</i>	
Nikita Moghe, Evgeniia Razumovskaia, Liane K Guillou, Ivan Vulić, Anna Korhonen and Alexandra Birch	3732
<i>A Robust Information-Masking Approach for Domain Counterfactual Generation</i>	
Pengfei Hong, Rishabh Bhardwaj, Navonil Majumder, Somak Aditya and Soujanya Poria	3756
<i>Misleading Relation Classifiers by Substituting Words in Texts</i>	
Tian Jiang, Yunqi Liu, Yan Feng, Yuqing Li and Xiaohui Cui	3770
<i>Automatic Table Union Search with Tabular Representation Learning</i>	
Xuming Hu, Shen Wang, Xiao Qin, Chuan Lei, Zhengyuan Shen, Christos Faloutsos, Asterios Katsifodimos, George Karypis, Lijie Wen and Philip S. Yu	3786
<i>Bidirectional Transformer Reranker for Grammatical Error Correction</i>	
Ying Zhang, Hidetaka Kamigaito and Manabu Okumura	3801

<i>Not Enough Data to Pre-train Your Language Model? MT to the Rescue!</i> Gorka Urbizu, Iñaki San Vicente, Xabier Saralegi and Ander Corral	3826
<i>UMSE: Unified Multi-scenario Summarization Evaluation</i> Shen Gao, Zhitao Yao, Chongyang Tao, Xiuying Chen, Pengjie Ren, Zhaochun Ren and Zhumin CHEN	3837
<i>Maximum Entropy Loss, the Silver Bullet Targeting Backdoor Attacks in Pre-trained Language Models</i> Zhengxiao Liu, Bowen Shen, Zheng Lin, Fali Wang and Weiping Wang	3850
<i>Improving Named Entity Recognition via Bridge-based Domain Adaptation</i> Jingyun Xu, Changmeng Zheng, Yi Cai and Tat-Seng Chua	3869
<i>SANTA: Separate Strategies for Inaccurate and Incomplete Annotation Noise in Distantly-Supervised Named Entity Recognition</i> Shuzheng Si, Zefan Cai, Shuang Zeng, Guoqiang Feng, Jiaying Lin and Baobao Chang	3883
<i>The State of Profanity Obfuscation in Natural Language Processing Scientific Publications</i> Debora Nozza and Dirk Hovy	3897
<i>Teacher and Student Models of Offensive Language in Social Media</i> Tharindu Ranasinghe and Marcos Zampieri	3910
<i>A Simple Yet Strong Domain-Agnostic De-bias Method for Zero-Shot Sentiment Classification</i> Yang Zhao, Tetsuya Nasukawa, Masayasu Muraoka and Bishwaranjan Bhattacharjee	3923
<i>Balancing the Effect of Training Dataset Distribution of Multiple Styles for Multi-Style Text Transfer</i> Debarati Das, David Ma and Dongyeop Kang	3932
<i>A Benchmark on Extremely Weakly Supervised Text Classification: Reconcile Seed Matching and Prompting Approaches</i> Zihan Wang, Tianle Wang, Dheeraj Mekala and Jingbo Shang	3944
<i>Ambiguity Meets Uncertainty: Investigating Uncertainty Estimation for Word Sense Disambiguation</i> Zhu Liu and Ying Liu	3963
<i>Zemi: Learning Zero-Shot Semi-Parametric Language Models from Multiple Tasks</i> Zhenhailong Wang, Xiaoman Pan, Dian Yu, Dong Yu, Jianshu Chen and Heng Ji	3978
<i>Why Can GPT Learn In-Context? Language Models Secretly Perform Gradient Descent as Meta-Optimizers</i> Damai Dai, Yutao Sun, Li Dong, Yaru Hao, Shuming Ma, Zhifang Sui and Furu Wei	4005
<i>Dramatic Conversation Disentanglement</i> Kent K. Chang, Danica Chen and David Bamman	4020
<i>Injecting Comparison Skills in Task-Oriented Dialogue Systems for Database Search Results Disambiguation</i> Yongil Kim, Yerin Hwang, Joongbo Shin, Hyunkyung Bae and Kyomin Jung	4047
<i>Emergent Modularity in Pre-trained Transformers</i> Zhengyan Zhang, Zhiyuan Zeng, Yankai Lin, Chaojun Xiao, Xiaozhi Wang, Xu Han, Zhiyuan Liu, Ruobing Xie, Maosong Sun and Jie Zhou	4066
<i>Universal Information Extraction with Meta-Pretrained Self-Retrieval</i> Xin Cong, Bowen Yu, Mengcheng Fang, Tingwen Liu, Haiyang Yu, Zhongkai Hu, Fei Huang, Yongbin Li and Bin Wang	4084

<i>SETI: Systematicity Evaluation of Textual Inference</i> Xiyan Fu and Anette Frank	4101
<i>Coarse-to-fine Few-shot Learning for Named Entity Recognition</i> Ruotian Ma, zhang lin, xuanting chen, Xin Zhou, Junzhe Wang, Tao Gui, Qi Zhang, Xiang Gao and Yun Wen Chen	4115
<i>Self-Evolution Learning for Discriminative Language Model Pretraining</i> Qihuang Zhong, Liang Ding, Juhua Liu, Bo Du and Dacheng Tao	4130
<i>QueryForm: A Simple Zero-shot Form Entity Query Framework</i> Zifeng Wang, Zizhao Zhang, Jacob Devlin, Chen-Yu Lee, Guolong Su, Hao Zhang, Jennifer Dy, Vincent Perot and Tomas Pfister	4146
<i>Search-Oriented Conversational Query Editing</i> Kelong Mao, Zhicheng Dou, Bang Liu, Hongjin Qian, Fengran Mo, Xiangli Wu, Xiaohua Cheng and Zhao Cao	4160
<i>TAPIR: Learning Adaptive Revision for Incremental Natural Language Understanding with a Two-Pass Model</i> Patrick Kahardipraja, Brielen Madureira and David Schlangen	4173
<i>Speaking the Language of Your Listener: Audience-Aware Adaptation via Plug-and-Play Theory of Mind</i> Ece Takmaz, Nicolo' Brandizzi, Mario Giulianelli, Sandro Pezzelle and Raquel Fernandez .	4198
<i>A Semi-Autoregressive Graph Generative Model for Dependency Graph Parsing</i> Ye Ma, Mingming Sun and Ping Li	4218
<i>AMR-TST: Abstract Meaning Representation-based Text Style Transfer</i> Kaize Shi, Xueyao Sun, LI HE, Dingxian Wang, Qing Li and Guandong Xu	4231
<i>Understanding the Cooking Process with English Recipe Text</i> YI FAN and Anthony Hunter	4244
<i>Follow the Wisdom of the Crowd: Effective Text Generation via Minimum Bayes Risk Decoding</i> Mirac Suzgun, Luke Melas-Kyriazi and Dan Jurafsky	4265
<i>RobustQA: Benchmarking the Robustness of Domain Adaptation for Open-Domain Question Answering</i> Rujun Han, Peng Qi, Yuhao Zhang, Lan Liu, Juliette Burger, William Yang Wang, zhiheng huang, Bing Xiang and Dan Roth	4294
<i>SenteCon: Leveraging Lexicons to Learn Human-Interpretable Language Representations</i> Victoria Lin and Louis-Philippe Morency	4312
<i>Reinforcement Learning for Topic Models</i> Jeremy J Costello and Marek Z Reformat	4332
<i>Contextualized Soft Prompts for Extraction of Event Arguments</i> Chien Van Nguyen, Hieu Man and Thien Huu Nguyen	4352
<i>TextVerifier: Robustness Verification for Textual Classifiers with Certifiable Guarantees</i> Siqi Sun and Wenjie Ruan	4362
<i>OASum: Large-Scale Open Domain Aspect-based Summarization</i> Xianjun Yang, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Xiaoman Pan, Linda R Petzold and Dong Yu	4381

<i>On the Limitations of Simulating Active Learning</i> Katerina Margatina and Nikolaos Aletras	4402
<i>Towards Alleviating the Object Bias in Prompt Tuning-based Factual Knowledge Extraction</i> Yuhang Wang, Dongyuan Lu, Chao Kong and Jitao Sang	4420
<i>vONTSS: vMF based semi-supervised neural topic modeling with optimal transport</i> Weijie Xu, Xiaoyu Jiang, Srinivasan Sengamedu Hanumantha Rao, Francis Iannacci and Jinjin Zhao	4433
<i>Bias Beyond English: Counterfactual Tests for Bias in Sentiment Analysis in Four Languages</i> Seraphina Goldfarb-Tarrant, Adam Lopez, Roi Blanco and Diego Marcheggiani	4458
<i>Complementary Explanations for Effective In-Context Learning</i> Xi Ye, Srinivasan Iyer, Asli Celikyilmaz, Veselin Stoyanov, Greg Durrett and Ramakanth Pasunuru	4469
<i>MISMATCH: Fine-grained Evaluation of Machine-generated Text with Mismatch Error Types</i> Keerthiram Murugesan, Sarathkrishna Swaminathan, Soham Dan, Subhajit Chaudhury, Chulaka Gunasekara, Maxwell Crouse, Diwakar Mahajan, Ibrahim Abdelaziz, Achille Fokoue, Pavan Kapanipathi, Salim Roukos and Alexander Gray	4485
<i>RHO: Reducing Hallucination in Open-domain Dialogues with Knowledge Grounding</i> Ziwei Ji, Zihan Liu, Nayeon Lee, Tiezheng Yu, Bryan Wilie, Min Zeng and Pascale Fung ..	4504
<i>Transformer Language Models Handle Word Frequency in Prediction Head</i> Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi and Kentaro Inui	4523
<i>Prompted LLMs as Chatbot Modules for Long Open-domain Conversation</i> Gibbeum Lee, Volker Hartmann, Jongho Park, Dimitris Papailiopoulos and Kangwook Lee	4536
<i>Prompt to be Consistent is Better than Self-Consistent? Few-Shot and Zero-Shot Fact Verification with Pre-trained Language Models</i> Fengzhu Zeng and Wei Gao	4555
<i>Model Analysis & Evaluation for Ambiguous Question Answering</i> Konstantinos Papakostas and Irene Papadopoulou	4570
<i>Debiasing should be Good and Bad: Measuring the Consistency of Debiasing Techniques in Language Models</i> Robert A. Morabito, Jad Kabbara and Ali Emami	4581
<i>Critic-Guided Decoding for Controlled Text Generation</i> Minbeom Kim, Hwanhee Lee, Kang Min Yoo, Joonsuk Park, Hwaran Lee and Kyomin Jung	4598
<i>MedNgage: A Dataset for Understanding Engagement in Patient-Nurse Conversations</i> Yan Wang, Heidi A.S. Donovan, Sabit Hassan and Malihe Alikhani	4613
<i>SEAG: Structure-Aware Event Causality Generation</i> Zhengwei Tao, Zhi Jin, Xiaoying Bai, Haiyan Zhao, Chengfeng Dou, yongqiang zhao, Fang Wang and Chongyang Tao	4631
<i>Large Language Models Can be Lazy Learners: Analyze Shortcuts in In-Context Learning</i> Ruixiang Tang, Dehan Kong, Longtao Huang and hui xue	4645
<i>A Two-Stage Decoder for Efficient ICD Coding</i> Thanh-Tung Nguyen, Viktor Schlegel, Abhinav Ramesh Kashyap and Stefan Winkler	4658

<i>Asymmetric feature interaction for interpreting model predictions</i>	
Xiaolei Lu, Jianghong Ma and Haode Zhang	4666
<i>Disagreement Matters: Preserving Label Diversity by Jointly Modeling Item and Annotator Label Distributions with DisCo</i>	
Tharindu Cyril Weerasooriya, Alexander Ororbia, Raj B Bhensadadia, Ashiqur KhudaBukhsh and Christopher Homan	4679
<i>Domain Aligned Prefix Averaging for Domain Generalization in Abstractive Summarization</i>	
Pranav Ajit Nair, Sukomal Pal and Pradeepika Verma	4696
<i>ClaimDiff: Comparing and Contrasting Claims on Contentious Issues</i>	
Miyoung Ko, Ingyu Seong, Hwaran Lee, Joonsuk Park, Minsuk Chang and Minjoon Seo ..	4711
<i>Unsupervised Paraphrasing of Multiword Expressions</i>	
Takashi Wada, Yuji Matsumoto, Timothy Baldwin and Jey Han Lau	4732
<i>G-Tuning: Improving Generalization of Pre-trained Language Models with Generative Adversarial Network</i>	
Rongxiang Weng, wen sen Cheng and Min Zhang	4747
<i>Unified Language Representation for Question Answering over Text, Tables, and Images</i>	
Bowen Yu, Cheng Fu, Haiyang Yu, Fei Huang and Yongbin Li	4756
<i>A Set Prediction Network For Extractive Summarization</i>	
Xiaoxia Cheng, Yongliang Shen and Weiming Lu	4766
<i>Geo-Seq2seq: Twitter User Geolocation on Noisy Data through Sequence to Sequence Learning</i>	
Jingyu Zhang, Alexandra DeLucia, Chenyu Zhang and Mark Dredze	4778
<i>Predicting Numerals in Text Using Nearest Neighbor Language Models</i>	
Taku Sakamoto and Akiko Aizawa	4795
<i>HonestBait: Forward References for Attractive but Faithful Headline Generation</i>	
Chih Yao Chen, Dennis Wu and Lun-Wei Ku	4810
<i>Few Shot Rationale Generation using Self-Training with Dual Teachers</i>	
Aditya Srikanth Veerubhotla, Lahari Poddar, Jun Yin, György Szarvas and Sharanya Eswaran	4825
<i>Towards Accurate Translation via Semantically Appropriate Application of Lexical Constraints</i>	
Yujin Baek, Koanho Lee, Dayeon Ki, Cheonbok Park, Hyoung-Gyu Lee and Jaegul Choo ..	4839
<i>NoisywikiHow: A Benchmark for Learning with Real-world Noisy Labels in Natural Language Processing</i>	
Tingting Wu, Xiao Ding, Minji Tang, Hao Zhang, Bing Qin and Ting Liu	4856
<i>Sampling Better Negatives for Distantly Supervised Named Entity Recognition</i>	
Lu Xu, Lidong Bing and Wei Lu	4874
<i>Prototype-Based Interpretability for Legal Citation Prediction</i>	
Chu Fei Luo, Rohan Bhambhoria, Samuel Dahan and Xiaodan Zhu	4883
<i>LMs stand their Ground: Investigating the Effect of Embodiment in Figurative Language Interpretation by Language Models</i>	
Philipp Wicke	4899

<i>Making Better Use of Training Corpus: Retrieval-based Aspect Sentiment Triplet Extraction via Label Interpolation</i>	
Guoxin Yu, Lemao Liu, Haiyun Jiang, Shuming Shi and Xiang Ao	4914
<i>Multi-Domain Dialogue State Tracking with Disentangled Domain-Slot Attention</i>	
Longfei Yang, Jiyi Li, Sheng Li and Takahiro Shinozaki	4928
<i>Improved Visual Story Generation with Adaptive Context Modeling</i>	
Zhangyin Feng, Yuchen Ren, Xinmiao Yu, Xiaocheng Feng, Duyu Tang, Shuming Shi and Bing Qin	4939
<i>Question-Interlocutor Scope Realized Graph Modeling over Key Utterances for Dialogue Reading Comprehension</i>	
Jiangnan Li, Mo Yu, Fandong Meng, Zheng Lin, Peng Fu, Weiping Wang and Jie Zhou	4956
<i>Speech-to-Speech Translation for a Real-world Unwritten Language</i>	
Peng-Jen Chen, Kevin Tran, Yilin Yang, Jingfei Du, Justine Kao, Yu-An Chung, Paden D Tomassello, Paul-Ambroise Augustin Duquenne, Holger Schwenk, Hongyu Gong, Hirofumi Inaguma, Sravya Popuri, Changhan Wang, Juan Pino, Wei-Ning Hsu and Ann Lee	4969
<i>Code Execution with Pre-trained Language Models</i>	
Chenxiao Liu, Shuai Lu, Weizhu Chen, Daxin Jiang, Alexey Svyatkovskiy, Shengyu Fu, Neel Sundaresan and Nan Duan	4984
<i>BertNet: Harvesting Knowledge Graphs with Arbitrary Relations from Pretrained Language Models</i>	
Shibo Hao, Bowen Tan, Kaiwen Tang, Bin Ni, Xiyan Shao, hengzhe zhang, Eric Xing and Zhiting Hu	5000
<i>Sequential Path Signature Networks for Personalised Longitudinal Language Modeling</i>	
Talia Tseriotou, Adam Tsakalidis, Peter Foster, Terence J Lyons and Maria Liakata	5016
<i>A Multi-modal Debiasing Model with Dynamical Constraint for Robust Visual Question Answering</i>	
Yu Li, Bojie Hu, Fengshuo Zhang, Yahan Yu, Jian Liu, Yufeng Chen and Jinan Xu	5032
<i>Trigger-Argument based Explanation for Event Detection</i>	
Yong Guan, Jiaoyan Chen, Freddy Lecue, Jeff Z. Pan, Juanzi Li and Ru Li	5046
<i>Interactive Concept Learning for Uncovering Latent Themes in Large Text Collections</i>	
Maria Leonor Pacheco, Tunazzina Islam, Lyle Ungar, Ming Yin and Dan Goldwasser	5059
<i>NormMark: A Weakly Supervised Markov Model for Socio-cultural Norm Discovery</i>	
Farhad Moghimifar, shilin qu, Tongtong Wu, Yuan-Fang Li and Gholamreza Haffari	5081
<i>VoteTRANS: Detecting Adversarial Text without Training by Voting on Hard Labels of Transformations</i>	
Hoang-Quoc Nguyen-Son, Seira Hidano, Kazuhide Fukushima, Shinsaku Kiyomoto and Isao Echizen	5090
<i>Fusion or Defusion? Flexible Vision-and-Language Pre-Training</i>	
Rongyi Sun, Ziran Li, Yifeng Ding, Qifan Wang, Jingang Wang, Haitao Zheng, Wei Wu and Yunsen Xian	5105
<i>COCKATIEL: Continuous Concept ranKed ATtribution with Interpretable ELEMents for explaining neural net classifiers on NLP</i>	
Fanny Jourdan, Agustin Martin Picard, Thomas Fel, Laurent Risser, Jean-Michel Loubes and Nicholas Asher	5120

<i>Code-Switched Text Synthesis in Unseen Language Pairs</i>	
I-Hung Hsu, Avik Ray, Shubham Garg, Nanyun Peng and Jing Huang	5137
<i>Imagination is All You Need! Curved Contrastive Learning for Abstract Sequence Modeling Utilized on Long Short-Term Dialogue Planning</i>	
Justus-Jonas Erker, Stefan Schaffer and Gerasimos Spanakis	5152
<i>Data-Efficient French Language Modeling with CamemBERTa</i>	
Wissam Antoun, Benoît Sagot and Djamé Seddah	5174
<i>Coupling Large Language Models with Logic Programming for Robust and General Reasoning from Text</i>	
Zhun Yang, Adam Ishay and Joohyung Lee	5186
<i>Evaluating the Factual Consistency of Large Language Models Through News Summarization</i>	
Derek Tam, Anisha Mascarenhas, Shiyue Zhang, Sarah Kwan, Mohit Bansal and Colin Raffel	5220
<i>Text Generation Model Enhanced with Semantic Information in Aspect Category Sentiment Analysis</i>	
Tu Dinh Tran, Kiyooki Shirai and Natthawut Kertkeidkachorn	5256
<i>Mind the Biases: Quantifying Cognitive Biases in Language Model Prompting</i>	
Ruixi Lin and Hwee Tou Ng	5269
<i>CodePrompt: Task-Agnostic Prefix Tuning for Program and Language Generation</i>	
YunSeok Choi and Jee-Hyong Lee	5282
<i>Honey, I Shrunk the Language: Language Model Behavior at Reduced Scale.</i>	
Vijeta Deshpande, Dan Pechi, Shree Thatte, Vladislav Lialin and Anna Rumshisky	5298
<i>Communication Efficient Federated Learning for Multilingual Neural Machine Translation with Adapter</i>	
Yi Liu, Xiaohan Bi, Lei Li, Sishuo Chen, Wenkai Yang and Xu Sun	5315
<i>Cross-task Knowledge Transfer for Extremely Weakly Supervised Text Classification</i>	
Seongmin Park, kyungho kim and Jihwa Lee	5329
<i>GVdoc - Graph-based Visual DOcument Classification</i>	
Fnu Mohbat, Mohammed J Zaki, Catherine Finegan-Dollak and Ashish Verma	5342
<i>A Sequence-to-Sequence&Set Model for Text-to-Table Generation</i>	
Tong Li, Zhihao Wang, Liangying Shao, Xuling Zheng, Xiaoli Wang and Jinsong Su	5358
<i>Automatic Readability Assessment for Closely Related Languages</i>	
Joseph Marvin Imperial and Ekaterina Kochmar	5371
<i>Towards Robust Ranker for Text Retrieval</i>	
Yucheng Zhou, Tao Shen, Xiubo Geng, Chongyang Tao, Can Xu, Guodong Long, Binxing Jiao and Daxin Jiang	5387
<i>Semi-Supervised Domain Adaptation for Emotion-Related Tasks</i>	
Mahshid Hosseini and Cornelia Caragea	5402
<i>Boosting Distress Support Dialogue Responses with Motivational Interviewing Strategy</i>	
Anuradha Welivita and Pearl Pu	5411

<i>ECOLA: Enhancing Temporal Knowledge Embeddings with Contextualized Language Representations</i> Zhen Han, Ruotong Liao, Jindong Gu, Yao Zhang, Zifeng Ding, Yujia Gu, Heinz Koepl, Hinrich Schütze and Volker Tresp	5433
<i>Gender-tuning: Empowering Fine-tuning for Debiasing Pre-trained Language Models</i> Somayeh Ghanbarzadeh, Yan Huang, Hamid Palangi, Radames Saul Cruz Moreno and Hamed Khanpour	5448
<i>TextObfuscator: Making Pre-trained Language Model a Privacy Protector via Obfuscating Word Representations</i> Xin Zhou, Yi Lu, Ruotian Ma, Tao Gui, Yuran Wang, yong ding, Yibo Zhang, Qi Zhang and Xuanjing Huang	5459
<i>Mini-Model Adaptation: Efficiently Extending Pretrained Models to New Languages via Aligned Shallow Training</i> Kelly Marchisio, Patrick Lewis, Yihong Chen and Mikel Artetxe	5474
<i>DSP: Discriminative Soft Prompts for Zero-Shot Entity and Relation Extraction</i> Bo Lv, Xin Liu, Shaojie Dai, Nayu Liu, Fan Yang, Ping Luo and Yue Yu	5491
<i>Exploring Robust Overfitting for Pre-trained Language Models</i> Bin Zhu and Yanghui Rao	5506
<i>Improving Cross-task Generalization of Unified Table-to-text Models with Compositional Task Configurations</i> Jifan Chen, Yuhao Zhang, Lan Liu, Rui Dong, Xinchu Chen, Patrick Ng, William Yang Wang and zhiheng huang	5523
<i>D-CALM: A Dynamic Clustering-based Active Learning Approach for Mitigating Bias</i> Sabit Hassan and Malihe Alikhani	5540
<i>Language Anisotropic Cross-Lingual Model Editing</i> Yang Xu, Yutai Hou, Wanxiang Che and Min Zhang	5554
<i>Diverse Retrieval-Augmented In-Context Learning for Dialogue State Tracking</i> Brendan King and Jeffrey Flanigan	5570
<i>Pre-Trained Language-Meaning Models for Multilingual Parsing and Generation</i> Chunliu Wang, Huiyuan Lai, Malvina Nissim and Johan Bos	5586
<i>Multi-modal Sarcasm Generation: Dataset and Solution</i> Wenye Zhao, Qingbao Huang, Dongsheng Xu and Peizhi Zhao	5601
<i>Rethinking Semi-supervised Learning with Language Models</i> Zhengxiang Shi, Francesco Tonolini, Nikolaos Aletras, Emine Yilmaz, Gabriella Kazai and Yunlong Jiao	5614
<i>Retrieval-Based Transformer for Table Augmentation</i> Michael Glass, Xueqing Wu, Ankita Rajaram Naik, Gaetano Rossiello and Alfio Gliozzo ..	5635
<i>ECG-QALM: Entity-Controlled Synthetic Text Generation using Contextual Q&A for NER</i> Karan Aggarwal, Henry Jin and Aitzaz Ahmad	5649
<i>Tokenization Impacts Multilingual Language Modeling: Assessing Vocabulary Allocation and Overlap Across Languages</i> Tomasz Limisiewicz, Jiří Balhar and David Mareček	5661

<i>The Whole Truth and Nothing But the Truth: Faithful and Controllable Dialogue Response Generation with Dataflow Transduction and Constrained Decoding</i>	
Hao Fang, Anusha Balakrishnan, Harsh Jhamtani, John Bufe, Jean Crawford, Jayant Krishnamurthy, Adam Pauls, Jason Eisner, Jacob Andreas and Dan Klein	5682
<i>Know What I don't Know: Handling Ambiguous and Unknown Questions for Text-to-SQL</i>	
Bing Wang, Yan Gao, Zhoujun Li and Jian-Guang LOU	5701
<i>Rethinking Document-Level Relation Extraction: A Reality Check</i>	
Jing Li, Yequan Wang, Shuai Zhang and Min Zhang	5715
<i>Optimizing Test-Time Query Representations for Dense Retrieval</i>	
Mujeen Sung, Jungsoo Park, Jaewoo Kang, Danqi Chen and Jinhyuk Lee	5731
<i>A Customized Text Sanitization Mechanism with Differential Privacy</i>	
Sai Chen, Fengran Mo, Yanhao Wang, Cen CHEN, Jian-Yun Nie, Chengyu Wang and Jamie Cui	5747
<i>LABO: Towards Learning Optimal Label Regularization via Bi-level Optimization</i>	
Peng Lu, Ahmad Rashid, Ivan Kobzyev, Mehdi Rezagholizadeh and Phillippe Langlais	5759
<i>Frustratingly Easy Label Projection for Cross-lingual Transfer</i>	
Yang Chen, Chao Jiang, Alan Ritter and Wei Xu	5775
<i>Enhancing Hierarchical Text Classification through Knowledge Graph Integration</i>	
Ye Liu, Kai Zhang, Zhenya Huang, Kehang Wang, Yanghai Zhang, Qi Liu and Enhong Chen	5797
<i>How Many Answers Should I Give? An Empirical Study of Multi-Answer Reading Comprehension</i>	
Chen Zhang, Jiuheg Lin, Xiao Liu, Yuxuan Lai, Yansong Feng and Dongyan Zhao	5811
<i>An Exploration of Encoder-Decoder Approaches to Multi-Label Classification for Legal and Biomedical Text</i>	
Yova Kementchedjhieva and Ilias Chalkidis	5828
<i>Domain Incremental Lifelong Learning in an Open World</i>	
Yi Dai, Hao Lang, Yinhe Zheng, Bowen Yu, Fei Huang and Yongbin Li	5844
<i>Improving Knowledge Graph Completion with Generative Hard Negative Mining</i>	
Zile Qiao, Wei Ye, Dingyao Yu, Tong Mo, Weiping Li and Shikun Zhang	5866
<i>Visually-Enhanced Phrase Understanding</i>	
Tsu-Yuan Hsu, Chen-An Li, Chao-Wei Huang and Yun-Nung Chen	5879
<i>Reasoning in Large Language Models Through Symbolic Math Word Problems</i>	
Vedant Gaur and Nikunj Saunshi	5889
<i>It's not Sexually Suggestive; It's Educational Separating Sex Education from Suggestive Content on TikTok videos</i>	
Enfa Rose George and Mihai Surdeanu	5904
<i>Dynamic Structured Neural Topic Model with Self-Attention Mechanism</i>	
Nozomu Miyamoto, Masaru Isonuma, Sho Takase, Junichiro Mori and Ichiro Sakata	5916
<i>Hybrid-Regressive Paradigm for Accurate and Speed-Robust Neural Machine Translation</i>	
Qiang Wang, Xinhui Hu and MING CHEN	5931
<i>Commonsense Knowledge Transfer for Pre-trained Language Models</i>	
Wangchunshu Zhou, Ronan Le Bras and Yejin Choi	5946

<i>Shielded Representations: Protecting Sensitive Attributes Through Iterative Gradient-Based Projection</i> Shadi Iskander, Kira Radinsky and Yonatan Belinkov	5961
<i>Focal Training and Tagger Decouple for Grammatical Error Correction</i> Minghuan Tan, Min Yang and Ruifeng Xu	5978
<i>LET: Leveraging Error Type Information for Grammatical Error Correction</i> Lingyu Yang, Hongjia Li, Lei Li, Chengyin Xu, Shutao Xia and Chun Yuan	5986
<i>On the Role of Parallel Data in Cross-lingual Transfer Learning</i> Machel Reid and Mikel Artetxe	5999
<i>CoMave: Contrastive Pre-training with Multi-scale Masking for Attribute Value Extraction</i> Xinnan Guo, Wentao Deng, Yongrui Chen, Yang Li, Mengdi Zhou, Guilin Qi, Tianxing Wu, Dong Yang, Liubin Wang and Yong Pan	6007
<i>Phrase Retrieval for Open Domain Conversational Question Answering with Conversational Dependency Modeling via Contrastive Learning</i> Soyeong Jeong, Jinheon Baek, Sung Ju Hwang and Jong Park	6019
<i>Unlearning Bias in Language Models by Partitioning Gradients</i> Charles Yu, Sullam Jeoung, Anish Kasi, Pengfei Yu and Heng Ji	6032
<i>Meta-training with Demonstration Retrieval for Efficient Few-shot Learning</i> Aaron Mueller, Kanika Narang, Lambert Mathias, Qifan Wang and Hamed Firooz	6049
<i>VCSUM: A Versatile Chinese Meeting Summarization Dataset</i> Han Wu, Mingjie Zhan, Haochen Tan, Zhaohui Hou, Ding Liang and Linqi Song	6065
<i>LEDA: a Large-Organization Email-Based Decision-Dialogue-Act Analysis Dataset</i> Mladen Karan, Prashant Khare, Ravi Shekhar, Stephen McQuistin, Ignacio Castro, Gareth Tyson, Colin Perkins, Patrick G.T. Healey and Matthew Purver	6080
<i>Negation Scope Refinement via Boundary Shift Loss</i> Yin Wu and Aixin Sun	6090
<i>Towards Diverse and Effective Question-Answer Pair Generation from Children Storybooks</i> Sugyeong Eo, Hyeonseok Moon, Jinsung Kim, Yuna Hur, Jeongwook Kim, SongEun Lee, Changwoo Chun, SUNGSOO PARK and Heuseok Lim	6100
<i>Pulling Out All The Full Stops: Punctuation Sensitivity in Neural Machine Translation and Evaluation</i> Prathyusha Jwalapuram	6116
<i>Reimagining Retrieval Augmented Language Models for Answering Queries</i> Wang-Chiew Tan, Yuliang Li, Pedro Rodriguez, Richard James, Xi Victoria Lin, Alon Halevy and Wen-tau Yih	6131
<i>Numeric Magnitude Comparison Effects in Large Language Models</i> Raj Sanjay Shah, Vijay Marupudi, Reba Koenen, Khushi Bhardwaj and Sashank Varma	6147
<i>Multi-Relational Probabilistic Event Representation Learning via Projected Gaussian Embedding</i> Linhai Zhang, Congzhi Zhang and Deyu Zhou	6162
<i>PragmaticCQA: A Dataset for Pragmatic Question Answering in Conversations</i> Peng Qi, Nina Du, Christopher D. Manning and Jing Huang	6175

<i>Modular and On-demand Bias Mitigation with Attribute-Removal Subnetworks</i> Lukas Hauzenberger, Shahed Masoudian, Deepak Kumar, Markus Schedl and Navid Rekasaz	6192
<i>Scientific Fact-Checking: A Survey of Resources and Approaches</i> Juraj Vladika and Florian Matthes	6215
<i>Uni-Encoder: A Fast and Accurate Response Selection Paradigm for Generation-Based Dialogue Systems</i> Chiyu Song, Hongliang He, Haofei Yu, Pengfei Fang, Leyang Cui and Zhenzhong Lan	6231
<i>DLAMA: A Framework for Curating Culturally Diverse Facts for Probing the Knowledge of Pretrained Language Models</i> Amr Keleg and Walid Magdy	6245
<i>Self-adaptive Context and Modal-interaction Modeling For Multimodal Emotion Recognition</i> Haozhe Yang, Xianqiang Gao, Jianlong Wu, Tian Gan, Ning Ding, Feijun Jiang and Liqiang Nie	6267
<i>Structure-Discourse Hierarchical Graph for Conditional Question Answering on Long Documents</i> Haowei Du, Yansong Feng, Chen Li, Yang Li, Yunshi Lan and Dongyan Zhao	6282
<i>COBRA Frames: Contextual Reasoning about Effects and Harms of Offensive Statements</i> Xuhui Zhou, Hao Zhu, Akhila Yerukola, Thomas Davidson, Jena D. Hwang, Swabha Swayamdipta and Maarten Sap	6294
<i>Distilling Calibrated Knowledge for Stance Detection</i> Yingjie Li and Cornelia Caragea	6316
<i>PTCSpell: Pre-trained Corrector Based on Character Shape and Pinyin for Chinese Spelling Correction</i> Xiao Wei, Jianbao Huang, Hang Yu and Qian Liu	6330
<i>Disentangling Text Representation With Counter-Template For Unsupervised Opinion Summarization</i> Yanyue Zhang and Deyu Zhou	6344
<i>Evaluation of Question Generation Needs More References</i> Shinhyeok Oh, Hyojun Go, Hyeongdon Moon, Yunsung Lee, Myeongho Jeong, Hyun Seung Lee and Seungtaek Choi	6358
<i>XtremeCLIP: Extremely Parameter-efficient Tuning for Low-resource Vision Language Understanding</i> Moming Tang, Chengyu Wang, Jianing Wang, Chuanqi Tan, Songfang Huang, Cen CHEN and Weining Qian	6368
<i>FACTUAL: A Benchmark for Faithful and Consistent Textual Scene Graph Parsing</i> Zhuang Li, Yuyang Chai, Terry Yue Zhuo, Lizhen Qu, Gholamreza Haffari, Fei Li, Donghong Ji and Quan Hung Tran	6377
<i>Target-Oriented Relation Alignment for Cross-Lingual Stance Detection</i> Ruikang Zhang, Nan Xu, Hanxuan Yang, Yuan Tian and Wenji Mao	6391
<i>NonFactS: NonFactual Summary Generation for Factuality Evaluation in Document Summarization</i> Amir Soleimani, Christof Monz and marcel worring	6405
<i>When to Read Documents or QA History: On Unified and Selective Open-domain QA</i> Kyungjae Lee, Sang-eun Han, Seung-won Hwang and Moontae Lee	6420

<i>Interpretable Automatic Fine-grained Inconsistency Detection in Text Summarization</i>	
Hou Pong Chan, Qi Zeng and Heng Ji	6433
<i>A Multi-dimensional study on Bias in Vision-Language models</i>	
Gabriele Ruggeri and Debora Nozza	6445
<i>Correction of Errors in Preference Ratings from Automated Metrics for Text Generation</i>	
Jan Deriu, Pius von Däniken, Don Tuggener and Mark Cieliebak	6456
<i>PEER: Pre-training ELECTRA Extended by Ranking</i>	
Ru He, Wei Wang, Songfang Huang and Fei Huang	6475
<i>ML-LMCL: Mutual Learning and Large-Margin Contrastive Learning for Improving ASR Robustness in Spoken Language Understanding</i>	
Xuxin Cheng, Bowen Cao, Qichen YE, Zhihong Zhu, Hongxiang Li and Yuexian Zou	6492
<i>Guiding Dialogue Agents to Complex Semantic Targets by Dynamically Completing Knowledge Graph</i>	
Yue Tan, Bo Wang, Anqi Liu, Dongming Zhao, Kun Huang, Ruifang He and Yuexian Hou ..	6506
<i>Chain of Thought Prompting Elicits Knowledge Augmentation</i>	
Dingjun Wu, Jing Zhang and Xinmei Huang	6519
<i>TACR: A Table Alignment-based Cell Selection Method for HybridQA</i>	
Jian Wu, YICHENG XU, Yan Gao, Jian-Guang LOU, Börje F. Karlsson and Manabu Okumura	6535
<i>Modeling Cross-Cultural Pragmatic Inference with Codenames Duet</i>	
Omar Shaikh, Caleb Ziems, William Held, Aryan J Pariani, Fred Morstatter and Diyi Yang ..	6550
<i>Werewolf Among Us: Multimodal Resources for Modeling Persuasion Behaviors in Social Deduction Games</i>	
Bolin Lai, Hongxin Zhang, Miao Liu, Aryan J Pariani, Fiona Ryan, Wenqi Jia, Shirley Anugrah Hayati, James M Rehg and Diyi Yang	6570
<i>Long to reign over us: A Case Study of Machine Translation and a New Monarch</i>	
Rebecca Knowles and Samuel Larkin	6589
<i>A Unified Generative Approach to Product Attribute-Value Identification</i>	
Keiji Shinzato, Naoki Yoshinaga, Yandi Xia and Wei-Te Chen	6599
<i>K-UniMorph: Korean Universal Morphology and its Feature Schema</i>	
Eunkyul Leah Jo, Kim Kyuwon, Xihan Wu, KyungTae Lim, Jungyeul Park and Chulwoo Park	6613
<i>How does the brain process syntactic structure while listening?</i>	
SUBBA REDDY OOTA, Mounika Marreddy, Manish Gupta and Raju Surampudi Bapi	6624
<i>Towards Imperceptible Document Manipulations against Neural Ranking Models</i>	
Xuanang Chen, Ben He, Zheng Ye, Le Sun and Yingfei Sun	6648
<i>Ask an Expert: Leveraging Language Models to Improve Strategic Reasoning in Goal-Oriented Dialogue Models</i>	
Qiang Zhang, Jason Naradowsky and Yusuke Miyao	6665
<i>SciReviewGen: A Large-scale Dataset for Automatic Literature Review Generation</i>	
Tetsu Kasanishi, Masaru Isonuma, Junichiro Mori and Ichiro Sakata	6695

<i>Revisiting Sample Size Determination in Natural Language Understanding</i> Ernie Chang, Muhammad Hassan Rashid, Pin-Jie Lin, Changsheng Zhao, Vera Demberg, Yangyang Shi and Vikas Chandra	6716
<i>TransESC: Smoothing Emotional Support Conversation via Turn-Level State Transition</i> Weixiang Zhao, Yanyan Zhao, Shilong Wang and Bing Qin	6725
<i>Residual Prompt Tuning: improving prompt tuning with residual reparameterization</i> Anastasiia Razdaibiedina, Yuning Mao, Madian Khabisa, Mike Lewis, Rui Hou, Jimmy Ba and Amjad Almahairi	6740
<i>Attend, Select and Eliminate: Accelerating Multi-turn Response Selection with Dual-attention-based Content Elimination</i> Jianxin Liang, Chang Liu, Chongyang Tao, Jiazhan Feng and Dongyan Zhao	6758
<i>Medical Dialogue Generation via Dual Flow Modeling</i> Kaishuai Xu, Wenjun Hou, Yi Cheng, Jian Wang and Wenjie Li	6771
<i>Listen, Decipher and Sign: Toward Unsupervised Speech-to-Sign Language Recognition</i> Liming Wang, Junrui Ni, Heting Gao, Jialu Li, Kai Chieh Chang, Xulin Fan, Junkai Wu, Mark Hasegawa-Johnson and Chang D. Yoo	6785
<i>Distinguishing Address vs. Reference Mentions of Personal Names in Text</i> Vinodkumar Prabhakaran, Aida Mostafazadeh Davani, Melissa Ferguson and Stav Atir	6801
<i>Low-ResourceText Classification: A Parameter-Free Classification Method with Compressors</i> Zhiying Jiang, Matthew Yang, Mikhail Tsirlin, Raphael Tang, Yiqin Dai and Jimmy Lin ...	6810
<i>LR-Sum: Summarization for Less-Resourced Languages</i> Chester Palen-Michel and Constantine Lignos	6829
<i>RQUGE: Reference-Free Metric for Evaluating Question Generation by Answering the Question</i> Alireza Mohammadshahi, Thomas Scialom, Majid Yazdani, Pouya Yanki, Angela Fan, James Henderson and Marzieh Saeidi	6845
<i>Unsupervised Semantic Variation Prediction using the Distribution of Sibling Embeddings</i> Taichi Aida and Danushka Bollegala	6868
<i>TransFormer: Slow-Fast Transformer for Machine Translation</i> Bei Li, Yi Jing, Xu Tan, Zhen Xing, Tong Xiao and Jingbo Zhu	6883
<i>Mitigating the Learning Bias towards Repetition by Self-Contrastive Training for Open-Ended Generation</i> Jian Guan and Minlie Huang	6897
<i>Digging out Discrimination Information from Generated Samples for Robust Visual Question Answering</i> Zhiquan Wen, Yaowei Wang, Mingkui Tan, Qingyao Wu and Qi Wu	6910
<i>Words as Gatekeepers: Measuring Discipline-specific Terms and Meanings in Scholarly Publications</i> Li Lucy, Jesse Dodge, David Bamman and Katherine A Keith	6929
<i>Trade-Offs Between Fairness and Privacy in Language Modeling</i> Cleo Matzken, Steffen Eger and Ivan Habernal	6948
<i>CSS: A Large-scale Cross-schema Chinese Text-to-SQL Medical Dataset</i> Hanchong Zhang, Jieyu Li, Lu Chen, Ruisheng Cao, Yunyan Zhang, Yu Huang, Yefeng Zheng and Kai Yu	6970

<i>Silver Syntax Pre-training for Cross-Domain Relation Extraction</i>	
Elisa Bassignana, Filip Ginter, Sampo Pyysalo, Rob van der Goot and Barbara Plank	6984
<i>FastDiff2: Revisiting and Incorporating GANs and Diffusion Models in High-Fidelity Speech Synthesis</i>	
Rongjie Huang, Yi Ren, Ziyue Jiang, Chenye Cui, Jinglin Liu and Zhou Zhao	6994
<i>Uncovering Hidden Consequences of Pre-training Objectives in Sequence-to-Sequence Models</i>	
Tannon Kew and Rico Sennrich	7010
<i>Exploring Anisotropy and Outliers in Multilingual Language Models for Cross-Lingual Semantic Sentence Similarity</i>	
Katharina Haemmerl, Alina Fastowski, Jindřich Libovický and Alexander Fraser	7023
<i>Revisiting Sentence Union Generation as a Testbed for Text Consolidation</i>	
Eran Hirsch, Valentina Pyatkin, Ruben Wolhandler, Avi Caciularu, Asi Shefer and Ido Dagan	7038
<i>Distilling Reasoning Capabilities into Smaller Language Models</i>	
Kumar Shridhar, Alessandro Stolfo and Mrinmaya Sachan	7059
<i>AlignSTS: Speech-to-Singing Conversion via Cross-Modal Alignment</i>	
Ruiqi Li, Rongjie Huang, Lichao Zhang, Jinglin Liu and Zhou Zhao	7074
<i>A New Task and Dataset on Detecting Attacks on Human Rights Defenders</i>	
Shihao Ran, Di Lu, Aoife Cahill, Joel Tetreault and Alejandro Jaimes	7089
<i>Improving Language Model Integration for Neural Machine Translation</i>	
Christian Herold, Yingbo Gao, Mohammad ZeinEldien and Hermann Ney	7114
<i>Type Enhanced BERT for Correcting NER Errors</i>	
Kuai Li, Chen Chen, Tao Yang, Tianming Du, Peijie Yu, dong du and Feng Zhang	7124
<i>Bridge the Gap Between CV and NLP! A Gradient-based Textual Adversarial Attack Framework</i>	
Lifan Yuan, YiChi Zhang, Yangyi Chen and Wei Wei	7132
<i>DUB: Discrete Unit Back-translation for Speech Translation</i>	
Dong Zhang, Rong Ye, Tom Ko, Mingxuan Wang and Yaqian Zhou	7147
<i>Knowledge Graph Embeddings using Neural Ito Process: From Multiple Walks to Stochastic Trajectories</i>	
Mojtaba Nayeri, Bo Xiong, Majid Mohammadi, Mst. Mahfuja Akter, Mirza Mohtashim Alam, Jens Lehmann and Steffen Staab	7165
<i>Leveraging Denoised Abstract Meaning Representation for Grammatical Error Correction</i>	
Hejing Cao and Dongyan Zhao	7180
<i>Prediction and Calibration: Complex Reasoning over Knowledge Graph with Bi-directional Directed Acyclic Graph Neural Network</i>	
Yao Xu, Shizhu He, Li Cai, Kang Liu and Jun Zhao	7189
<i>Prompt-Based Metric Learning for Few-Shot NER</i>	
Yanru Chen, Yanan Zheng and Zhilin Yang	7199
<i>OpenPI-C: A Better Benchmark and Stronger Baseline for Open-Vocabulary State Tracking</i>	
Xueqing Wu, Sha Li and Heng Ji	7213
<i>I run as fast as a rabbit, can you? A Multilingual Simile Dialogues Datasets</i>	
Longxuan Ma, Wei-Nan Zhang, Shuhan Zhou, churui sun, Changxin Ke and Ting Liu	7223

<i>Controllable Conversation Generation with Conversation Structures via Diffusion Models</i>	
Jiaao Chen and Diyi Yang	7238
<i>Few-shot Low-resource Knowledge Graph Completion with Reinforced Task Generation</i>	
Shichao Pei, qiannan zhang and Xiangliang Zhang	7252
<i>Incomplete Utterance Rewriting as Sequential Greedy Tagging</i>	
Yunshan Chen	7265
<i>Exploiting Commonsense Knowledge about Objects for Visual Activity Recognition</i>	
Tianyu Jiang and Ellen Riloff	7277
<i>Tucker Decomposition with Frequency Attention for Temporal Knowledge Graph Completion</i>	
Likang Xiao, Richong Zhang, Zijie Chen and Junfan Chen	7286
<i>Another Dead End for Morphological Tags? Perturbed Inputs and Parsing</i>	
Alberto Muñoz-Ortiz and David Vilares	7301
<i>HeGeL: A Novel Dataset for Geo-Location from Hebrew Text</i>	
Tzuf Paz-Argaman, Tal Bauman, Itai Mondshine, Itzhak Omer, Sagi Dalyot and Reut Tsarfaty	
7311	
<i>Modeling Adversarial Attack on Pre-trained Language Models as Sequential Decision Making</i>	
Xuanjie Fang, Sijie Cheng, Yang Liu and wei wang	7322
<i>Towards Robust Personalized Dialogue Generation via Order-Insensitive Representation Regularization</i>	
Liang CHEN, Hongru Wang, Yang Deng, Wai Chung Kwan, Zezhong WANG and Kam-Fai Wong	
7337	
<i>Cost-effective Distillation of Large Language Models</i>	
Sayantan Dasgupta, Trevor Cohn and Timothy Baldwin	7346
<i>Task-Optimized Adapters for an End-to-End Task-Oriented Dialogue System</i>	
Namo Bang, Jeehyun Lee and Myoung-Wan Koo	7355
<i>I Spy a Metaphor: Large Language Models and Diffusion Models Co-Creat Visual Metaphors</i>	
Tuhin Chakrabarty, Arkadiy Saakyan, Olivia Winn, Artemis Panagopoulou, Yue Yang, Marianna	
Apidianaki and Smaranda Muresan	7370
<i>Text Augmentation Using Dataset Reconstruction for Low-Resource Classification</i>	
Adir Rahamim, Guy Uziel, Esther Goldbraich and Ateret Anaby Tavor	7389
<i>LaSQe: Improved Zero-Shot Classification from Explanations Through Quantifier Modeling and Curriculum Learning</i>	
Sayan Ghosh, Rakesh R. Menon and Shashank Srivastava	7403
<i>Learned Adapters Are Better Than Manually Designed Adapters</i>	
Yuming Zhang, Peng Wang, Ming Tan and Wei Zhu	7420
<i>Automatic Identification of Code-Switching Functions in Speech Transcripts</i>	
Ritu Madhura Belani and Jeffrey Flanigan	7438
<i>Federated Domain Adaptation for Named Entity Recognition via Distilling with Heterogeneous Tag Sets</i>	
Rui Wang, Tong Yu, Junda Wu, Handong Zhao, Sungchul Kim, Ruiyi Zhang, Subrata Mitra and	
Ricardo Henao	7449

<i>Interpreting Sentiment Composition with Latent Semantic Tree</i> Zhongtao Jiang, Yuanzhe Zhang, Cao Liu, Jiansong Chen, Jun Zhao and Kang Liu	7464
<i>Beyond Positive Scaling: How Negation Impacts Scaling Trends of Language Models</i> Yuhui Zhang, Michihiro Yasunaga, Zhengping Zhou, Jeff Z. HaoChen, James Zou, Percy Liang and Serena Yeung	7479
<i>Contrastive Training Improves Zero-Shot Classification of Semi-structured Documents</i> Muhammad Khalifa, Yogarshi Vyas, Shuai Wang, Graham Horwood, Sunil Mallya and Miguel Ballesteros	7499
<i>Extracting Shopping Interest-Related Product Types from the Web</i> Yinghao Li, Colin Lockard, Prashant Shiralkar and Chao Zhang	7509
<i>Multilingual Pre-training with Self-supervision from Global Co-occurrence Information</i> Xi Ai and Bin Fang	7526
<i>Low-Rank Updates of pre-trained Weights for Multi-Task Learning</i> Alexandre Daniel AUDIBERT, Massih R Amini, Konstantin Usevich and Marianne Clausel	7544
<i>Sequential Integrated Gradients: a simple but effective method for explaining language models</i> Joseph Enguehard	7555
<i>DiffuDetox: A Mixed Diffusion Model for Text Detoxification</i> Griffin F Floto, Mohammad Mahdi Abdollah Pour, Parsa Farinneya, Zhenwei Tang, Ali Pesaran- ghader, Manasa Bharadwaj and Scott P Sanner	7566
<i>Separating Context and Pattern: Learning Disentangled Sentence Representations for Low-Resource Extractive Summarization</i> Ruifeng Yuan, Shichao Sun, zili Wang, Ziqiang Cao and Wenjie Li	7575
<i>Disentangling Reasoning Capabilities from Language Models with Compositional Reasoning Transfor- mers</i> Wanjun Zhong, Tingting Ma, Jiahai Wang, Jian Yin, Tiejun Zhao, Chin-Yew Lin and Nan Duan	7587
<i>Towards Argument-Aware Abstractive Summarization of Long Legal Opinions with Summary Reranking</i> Mohamed Elaraby, Yang Zhong and Diane Litman	7601
<i>Probabilistic Transformer: A Probabilistic Dependency Model for Contextual Word Representation</i> Haoyi Wu and Kewei Tu	7613
<i>Joint Speech Transcription and Translation: Pseudo-Labeling with Out-of-Distribution Data</i> Mozhdeh Gheini, Tatiana Likhomanenko, Matthias Sperber and Hendra Setiawan	7637
<i>Word-level Prefix/Suffix Sense Detection: A Case Study on Negation Sense with Few-shot Learning</i> Yameng Li, Zicheng Li, Ying Chen and Shoushan Li	7651
<i>End-to-End Simultaneous Speech Translation with Differentiable Segmentation</i> Shaolei Zhang and Yang Feng	7659
<i>Joint Generator-Ranker Learning for Natural Language Generation</i> Weizhou Shen, Yeyun Gong, Yelong Shen, Song Wang, Xiaojun Quan, Nan Duan and Weizhu Chen	7681
<i>Multilingual Sequence-to-Sequence Models for Hebrew NLP</i> Matan Eyal, Hila Noga, Roei Aharoni, Idan Szpektor and Reut Tsarfaty	7700

<i>Multilingual Knowledge Graph Completion from Pretrained Language Models with Knowledge Constraints</i>	Ran Song, Shizhu He, Shengxiang Gao, Li Cai, Kang Liu, Zhengtao YU and Jun Zhao	7709
<i>Towards Better Hierarchical Text Classification with Data Generation</i>	Yue Wang, Dan Qiao, Juntao Li, Jinxiong Chang, Qishen Zhang, Zhongyi Liu, Guannan Zhang and Min Zhang	7722
<i>History repeats: Overcoming catastrophic forgetting for event-centric temporal knowledge graph completion</i>	Mehrnoosh Mirtaheri, Mohammad Rostami and Aram Galstyan	7740
<i>Multi-Agent Language Learning: Symbolic Mapping</i>	Yicheng Feng and Zongqing Lu	7756
<i>Scaling Laws for BERT in Low-Resource Settings</i>	Gorka Urbizu, Iñaki San Vicente, Xabier Saralegi, Rodrigo Agerri and Aitor Soroa	7771
<i>Pre-trained Language Model with Prompts for Temporal Knowledge Graph Completion</i>	Wenjie Xu, Ben Liu, Miao Peng, Xu Jia and Min Peng	7790
<i>Is Continuous Prompt a Combination of Discrete Prompts? Towards a Novel View for Interpreting Continuous Prompts</i>	Tianjie Ju, Yubin Zheng, Hanyi Wang, Haodong Zhao and Gongshen Liu	7804
<i>Putting Natural in Natural Language Processing</i>	Grzegorz Chrupała	7820
<i>Impact of Adversarial Training on Robustness and Generalizability of Language Models</i>	Enes Altinisik, Hassan Sajjad, Husrev Taha Sencar, Safa Messaoud and Sanjay Chawla	7828
<i>Benchmarking Diverse-Modal Entity Linking with Generative Models</i>	Sijia Wang, Alexander Hanbo Li, Henghui Zhu, Sheng Zhang, Pramuditha Perera, Chung-Wei Hang, Jie Ma, William Yang Wang, Zhiguo Wang, Vittorio Castelli, Bing Xiang and Patrick Ng .	7841
<i>Improving Empathetic Dialogue Generation by Dynamically Infusing Commonsense Knowledge</i>	Hua Cai, Xuli Shen, Qing Xu, Weilin Shen, Xiaomei Wang, Weifeng Ge, Xiaoqing Zheng and Xiangyang Xue	7858
<i>Additive manifesto decomposition: A policy domain aware method for understanding party positioning</i>	Tanise Ceron, Dmitry Nikolaev and Sebastian Padó	7874
<i>Similarizing the Influence of Words with Contrastive Learning to Defend Word-level Adversarial Text Attack</i>	Pengwei Zhan, Jing Yang, He Wang, Chao Zheng, Xiao Huang and Liming Wang	7891
<i>Responsibility Perspective Transfer for Italian Femicide News</i>	Gosse Minnema, Huiyuan Lai, Benedetta Muscato and Malvina Nissim	7907
<i>Stereotypes and Smut: The (Mis)representation of Non-cisgender Identities by Text-to-Image Models</i>	Eddie L. Ungless, Bjorn Ross and Anne Lauscher	7919
<i>Fine-grained Artificial Neurons in Audio-transformers for Disentangling Neural Auditory Encoding</i>	Mengyue Zhou, Xu Liu, David Liu, Zihao Wu, Zhengliang Liu, Lin Zhao, Dajiang Zhu, Lei Guo, Junwei Han, Tianming Liu and Xintao Hu	7943
<i>Deeply Coupled Cross-Modal Prompt Learning</i>	Xuejing Liu, Wei Tang, Jinghui Lu, Rui Zhao, Zhaojun Guo and Fei Tan	7957

<i>Opinion Tree Parsing for Aspect-based Sentiment Analysis</i>	
Xiaoyi Bao, Xiaotong Jiang, Zhongqing Wang, Yue Zhang and Guodong Zhou	7971
<i>CoMix: Guide Transformers to Code-Mix using POS structure and Phonetics</i>	
Gaurav Arora, Srujana Merugu and Vivek Varadarajan Sembium	7985
<i>Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes</i>	
Cheng-Yu Hsieh, Chun-Liang Li, CHIH-KUAN YEH, Hootan Nakhost, Yasuhisa Fujii, Alex Jason Ratner, Ranjay Krishna, Chen-Yu Lee and Tomas Pfister	8003
<i>Prosody-TTS: Improving Prosody with Masked Autoencoder and Conditional Diffusion Model For Expressive Text-to-Speech</i>	
Rongjie Huang, Chunlei Zhang, Yi Ren, Zhou Zhao and Dong Yu	8018
<i>Duplex Diffusion Models Improve Speech-to-Speech Translation</i>	
Xianchao Wu	8035
<i>Global and Local Hierarchy-aware Contrastive Framework for Implicit Discourse Relation Recognition</i>	
Yuxin Jiang, Linhan Zhang and Wei Wang	8048
<i>PreQuant: A Task-agnostic Quantization Approach for Pre-trained Language Models</i>	
Zhuocheng Gong, Jiahao Liu, Qifan Wang, Yang Yang, Jingang Wang, Wei Wu, Yunsen Xian, Dongyan Zhao and Rui Yan	8065
<i>Synthetic Pre-Training Tasks for Neural Machine Translation</i>	
Zexue He, Graeme Blackwood, Rameswar Panda, Julian McAuley and Rogerio Feris	8080
<i>IDOL: Indicator-oriented Logic Pre-training for Logical Reasoning</i>	
Zihang Xu, Ziqing Yang, Yiming Cui and Shijin Wang	8099
<i>Adversarial Training for Low-Resource Disfluency Correction</i>	
Vineet Bhat, Preethi Jyothi and Pushpak Bhattacharyya	8112
<i>Computer says No": The Case Against Empathetic Conversational AI</i>	
Alba A Cercas Curry and Amanda Cercas Curry	8123
<i>Stubborn Lexical Bias in Data and Models</i>	
Sofia Serrano, Jesse Dodge and Noah A. Smith	8131
<i>Distilling Efficient Language-Specific Models for Cross-Lingual Transfer</i>	
Alan Ansell, Edoardo Maria Ponti, Anna Korhonen and Ivan Vulić	8147
<i>An Extensive Exploration of Back-Translation in 60 Languages</i>	
Paul McNamee and Kevin Duh	8166
<i>AoM: Detecting Aspect-oriented Information for Multimodal Aspect-Based Sentiment Analysis</i>	
Ru Zhou, Wenya Guo, Xumeng Liu, Shenglong Yu, Ying Zhang and Xiaojie Yuan	8184
<i>Forecasting Earnings Surprises from Conference Call Transcripts</i>	
Ross Koval, Nicholas Andrews and Xifeng Yan	8197
<i>MTCue: Learning Zero-Shot Control of Extra-Textual Attributes by Leveraging Unstructured Context in Neural Machine Translation</i>	
Sebastian T. Vincent, Robert James Flynn and Carolina Scarton	8210
<i>Evaluation for Change</i>	
Rishi Bommasani	8227

<i>Reconstruction Probing</i>	
Najoung Kim, Jatin Khilnani, Alex Warstadt and Abdelrahim Qaddoumi	8240
<i>Towards Distribution-shift Robust Text Classification of Emotional Content</i>	
Luana Bulla, Aldo Gangemi and Misael Mongiovi’	8256
<i>Multi-lingual and Multi-cultural Figurative Language Understanding</i>	
Anubha Kabra, Emmy Liu, Simran Khanuja, Alham Fikri Aji, Genta Indra Winata, Samuel Cahyawijaya, Anuoluwapo Aremu, Perez Ogayo and Graham Neubig	8269
<i>Open-WikiTable : Dataset for Open Domain Question Answering with Complex Reasoning over Table</i>	
Sunjun Kweon, Yeonsu Kwon, Seonhee Cho, Yohan Jo and Edward Choi	8285
<i>What In-Context Learning Learns In-Context: Disentangling Task Recognition and Task Learning</i>	
Jane Pan, Tianyu Gao, Howard Chen and Danqi Chen	8298
<i>Cross-Lingual Retrieval Augmented Prompt for Low-Resource Languages</i>	
Ercong Nie, Sheng Liang, Helmut Schmid and Hinrich Schütze	8320
<i>Unsupervised Summarization Re-ranking</i>	
Mathieu Ravaut, Shafiq Joty and Nancy Chen	8341
<i>GRACE: Gradient-guided Controllable Retrieval for Augmenting Attribute-based Text Generation</i>	
Zhihua Wen, Zhiliang Tian, Zhen Huang, Yuxin Yang, Zexin Jian, Changjian Wang and Dongsheng Li	8377
<i>So many design choices: Improving and interpreting neural agent communication in signaling games</i>	
Timothee Bernard and Timothee Mickus	8399
<i>Constructing Word-Context-Coupled Space Aligned with Associative Knowledge Relations for Interpretable Language Modeling</i>	
Fanyu Wang and Zhenping Xie	8414
<i>Fixed Input Parameterization for Efficient Prompting</i>	
Eunbi Choi, Yongrae Jo, Joel Jang, Joonwon Jang and Minjoon Seo	8428
<i>Data Augmentation for Low-Resource Keyphrase Generation</i>	
Krishna K Garg, Jishnu Ray Chowdhury and Cornelia Caragea	8442
<i>BigVideo: A Large-scale Video Subtitle Translation Dataset for Multimodal Machine Translation</i>	
Liyan Kang, Luyang Huang, Ningxin Peng, Peihao Zhu, Zewei Sun, Shanbo Cheng, Mingxuan Wang, Degen Huang and Jinsong Su	8456
<i>Constructing Procedural Graphs with Multiple Dependency Relations: A New Dataset and Baseline</i>	
Haopeng Ren, Yushi Zeng, Yi Cai, Bihan Zhou and Zetao Lian	8474
<i>Multi-Dimensional Evaluation of Text Summarization with In-Context Learning</i>	
Sameer Jain, Vaishakh Keshava, Swarnashree Mysore Sathyendra, Patrick Fernandes, Pengfei Liu, Graham Neubig and Chunting Zhou	8487
<i>Learning to Rank Utterances for Query-Focused Meeting Summarization</i>	
Xingxian Liu and Yajing Xu	8496
<i>Neural Architecture Search for Parameter-Efficient Fine-tuning of Large Pre-trained Language Models</i>	
Neal G Lawton, Anoop Kumar, Govind Thattai, Aram Galstyan and Greg Ver Steeg	8506

<i>Aligning Offline Metrics and Human Judgments of Value for Code Generation Models</i>	
Victor Dibia, Adam Fourney, Gagan Bansal, Forough Poursabzi-Sangdeh, Han Liu and Saleema Amershi	8516
<i>Do transformer models do phonology like a linguist?</i>	
Saliha Muradoglu and Mans Hulden	8529
<i>DiMS: Distilling Multiple Steps of Iterative Non-Autoregressive Transformers for Machine Translation</i>	
Sajad Norouzi, Rasa Hosseinzadeh, Felipe Perez and Maksims Volkovs	8538
<i>Retrieval-augmented Video Encoding for Instructional Captioning</i>	
Yeonjoon Jung, Minsoo Kim, Seungtaek Choi, Jihyuk Kim, Minji Seo and Seung-won Hwang	8554
<i>Bi-level Finetuning with Task-dependent Similarity Structure for Low-resource Training</i>	
Sai Ashish Somayajula, Lifeng Jin, Linfeng Song, Haitao Mi and Dong Yu	8569
<i>Kanbun-LM: Reading and Translating Classical Chinese in Japanese Methods by Language Models</i>	
Hao Wang, Hirofumi Shimizu and Daisuke Kawahara	8589
<i>Adaptive Attention for Sparse-based Long-sequence Transformer</i>	
Xuanyu Zhang, Zhepeng Lv and Qing Yang	8602
<i>Sentiment Analysis using the Relationship between Users and Products</i>	
Natthawut Kertkeidkachorn and Kiyoaki Shirai	8611
<i>Entropy-guided Vocabulary Augmentation of Multilingual Language Models for Low-resource Tasks</i>	
Arijit Nag, bidisha Samanta, Animesh Mukherjee, Niloy Ganguly and Soumen Chakrabarti	8619
<i>Class-Adaptive Self-Training for Relation Extraction with Incompletely Annotated Training Data</i>	
Qingyu Tan, Lu Xu, Lidong Bing and Hwee Tou Ng	8630
<i>Solving Cosine Similarity Underestimation between High Frequency Words by ℓ_2 Norm Discounting</i>	
Saeth Wannasuphprasit, Yi Zhou and Danushka Bollegala	8644
<i>Do Large Language Models Know What They Don't Know?</i>	
Zhangyue Yin, Qiushi Sun, Qipeng Guo, Jiawen Wu, Xipeng Qiu and Xuanjing Huang . . .	8653
<i>AltCLIP: Altering the Language Encoder in CLIP for Extended Language Capabilities</i>	
Zhongzhi Chen, Guang Liu, Bo-Wen Zhang, Qinghong Yang and Ledell Y. Wu	8666
<i>RHGN: Relation-gated Heterogeneous Graph Network for Entity Alignment in Knowledge Graphs</i>	
Xukai Liu, Kai Zhang, Ye Liu, Enhong Chen, Zhenya Huang, Linan Yue and Jiaxian Yan . .	8683
<i>Feature Interactions Reveal Linguistic Structure in Language Models</i>	
Jaap Jumelet and Willem Zuidema	8697
<i>Clustering-Aware Negative Sampling for Unsupervised Sentence Representation</i>	
Jinghao Deng, Fanqi Wan, Tao Yang, Xiaojun Quan and Rui Wang	8713
<i>An Effective Deployment of Contrastive Learning in Multi-label Text Classification</i>	
Nankai Lin, Guanqiu Qin, Gang Ji Wang, Dong Zhou and Aimin Yang	8730
<i>Segment-Level and Category-Oriented Network for Knowledge-Based Referring Expression Comprehension</i>	
Yuqi Bu, Xin Wu, Liuwu Li, Yi Cai, Qiong Liu and Qingbao Huang	8745

<i>MVP: Multi-task Supervised Pre-training for Natural Language Generation</i> Tianyi Tang, Junyi Li, Wayne Xin Zhao and Ji-Rong Wen	8758
<i>From Alignment to Entailment: A Unified Textual Entailment Framework for Entity Alignment</i> Yu Zhao, Yike Wu, Xiangrui Cai, Ying Zhang, Haiwei Zhang and Xiaojie Yuan	8795
<i>It is a Bird Therefore it is a Robin: On BERT's Internal Consistency Between Hypernym Knowledge and Logical Words</i> Nicolas Pierre Guerin and Emmanuel Chemla	8807
<i>Defending against Insertion-based Textual Backdoor Attacks via Attribution</i> Jiazhao Li, Zhuofeng Wu, Wei Ping, Chaowei Xiao and V.G.Vinod Vydiswaran	8818
<i>ActiveAED: A Human in the Loop Improves Annotation Error Detection</i> Leon Weber and Barbara Plank	8834
<i>Assessing Word Importance Using Models Trained for Semantic Tasks</i> David Javorsky, Ondřej Bojar and François Yvon	8846
<i>In-context Examples Selection for Machine Translation</i> Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer and Marjan Ghazvininejad	8857
<i>PropSegmEnt: A Large-Scale Corpus for Proposition-Level Segmentation and Entailment Recognition</i> Sihao Chen, Senaka Buthpitiya, Alex Fabrikant, Dan Roth and Tal Schuster	8874
<i>CIF-PT: Bridging Speech and Text Representations for Spoken Language Understanding via Continuous Integrate-and-Fire Pre-Training</i> Linhao Dong, Zhecheng An, Peihao Wu, Jun Zhang, Lu Lu and MA Zejun	8894
<i>Improving Diachronic Word Sense Induction with a Nonparametric Bayesian method</i> Ashjan Alsulaimani and Erwan Moreau	8908
<i>What to Fuse and How to Fuse: Exploring Emotion and Personality Fusion Strategies for Explainable Mental Disorder Detection</i> Sourabh Satish Zanwar, Xiaofei Li, Daniel Wiechmann, Yu Qiao and Elma Kerz	8926
<i>Adaptive Contrastive Knowledge Distillation for BERT Compression</i> Jinyang Guo, Jiaheng Liu, Zining Wang, Yuqing Ma, Ruihao Gong, Ke Xu and Xianglong Liu	8941
<i>Fourier Transformer: Fast Long Range Modeling by Removing Sequence Redundancy with FFT Operator</i> Ziwei He, Meng Yang, Minwei Feng, Jingcheng Yin, Xinbing Wang, Jingwen Leng and Zhouhan Lin	8954
<i>Zero-Shot Classification by Logical Reasoning on Natural Language Explanations</i> Chi Han, Hengzhi Pei, Xinya Du and Heng Ji	8967
<i>Dual-Gated Fusion with Prefix-Tuning for Multi-Modal Relation Extraction</i> Qian Li, Shu Guo, Cheng Ji, Xutan Peng, Shiyao Cui, Jianxin Li and Lihong Wang	8982
<i>Pruning Pre-trained Language Models with Principled Importance and Self-regularization</i> Siyu Ren and Kenny Q. Zhu	8995
<i>The Magic of IF: Investigating Causal Reasoning Abilities in Large Language Models of Code</i> Xiao Liu, Da Yin, Chen Zhang, Yansong Feng and Dongyan Zhao	9009

<i>Learning to Leverage High-Order Medical Knowledge Graph for Joint Entity and Relation Extraction</i> Zhe YANG, Yi Huang and Junlan Feng	9023
<i>Data-Efficient Finetuning Using Cross-Task Nearest Neighbors</i> Hamish Ivison, Noah A. Smith, Hannaneh Hajishirzi and Pradeep Dasigi	9036
<i>CoAug: Combining Augmentation of Labels and Labelling Rules</i> Rakesh R. Menon, Bingqing Wang, Jun Araki, Zhengyu Zhou, Zhe Feng and Liu Ren	9062
<i>Entity-to-Text based Data Augmentation for various Named Entity Recognition Tasks</i> Xuming Hu, Yong Jiang, Aiwei Liu, Zhongqiang Huang, Pengjun Xie, Fei Huang, Lijie Wen and Philip S. Yu	9072
<i>World Models for Math Story Problems</i> Andreas Opedal, Niklas Stoehr, Abulhair Saparov and Mrinmaya Sachan	9088
<i>AutoMoE: Heterogeneous Mixture-of-Experts with Adaptive Computation for Efficient Neural Machine Translation</i> Ganesh Jawahar, Subhabrata Mukherjee, Xiaodong Liu, Young Jin Kim, Muhammad Abdul-Mageed, Laks Lakshmanan, V.S., Ahmed Hassan Awadallah, Sebastien Bubeck and Jianfeng Gao	9116
<i>Language Agnostic Multilingual Information Retrieval with Contrastive Learning</i> Xiyang Hu, Xinchu Chen, Peng Qi, Deguang Kong, Kunlun Liu, William Yang Wang and zhiheng huang	9133
<i>Easy to Decide, Hard to Agree: Reducing Disagreements Between Saliency Methods</i> Josip Jukić, Martin Tutek and Jan Snajder	9147
<i>Enhancing Cross-lingual Transfer via Phonemic Transcription Integration</i> Hoang Nguyen, Chenwei Zhang, Tao Zhang, Eugene Rohrbaugh and Philip S Yu	9163
<i>Human-in-the-loop Abstractive Dialogue Summarization</i> Jiaao Chen, Mohan Dodda and Diyi Yang	9176
<i>A Multi-task Learning Framework for Quality Estimation</i> Sourabh Dattatray Deoghare, Paramveer Choudhary, Diptesh Kanojia, Tharindu Ranasinghe, Puspak Bhattacharyya and Constantin Orăsan	9191
<i>The Devil is in the Details: On the Pitfalls of Event Extraction Evaluation</i> Hao Peng, Xiaozhi Wang, Feng Yao, Kaisheng Zeng, Lei Hou, Juanzi Li, Zhiyuan Liu and Weixing Shen	9206
<i>Yes, this Way! Learning to Ground Referring Expressions into Actions with Intra-episodic Feedback from Supportive Teachers</i> Philipp Sadler, Sherzod Hakimov and David Schlangen	9228
<i>Investigating Transformer-Guided Chaining for Interpretable Natural Logic Reasoning</i> Kanagasabai Rajaraman, Saravanan Rajamanickam and Wei Shi	9240
<i>Multilingual Multi-Figurative Language Detection</i> Huiyuan Lai, Antonio Toral and Malvina Nissim	9254
<i>Zero-shot Visual Question Answering with Language Model Feedback</i> Yifan Du, Junyi Li, Tianyi Tang, Wayne Xin Zhao and Ji-Rong Wen	9268
<i>Prompted Opinion Summarization with GPT-3.5</i> Adithya Bhaskar, Alex Fabbri and Greg Durrett	9282

<i>Sentence Ordering with a Coherence Verifier</i>	
Sainan Jia, Wei Song, Jiefu Gong, Shijin Wang and Ting Liu	9301
<i>GUMSum: Multi-Genre Data and Evaluation for English Abstractive Summarization</i>	
Yang Janet Liu and Amir Zeldes	9315
<i>Improving Grammatical Error Correction with Multimodal Feature Integration</i>	
Tao Fang, jinpeng hu, Derek F. Wong, Xiang Wan, Lidia S. Chao and Tsung-Hui Chang . . .	9328
<i>Teaching the Pre-trained Model to Generate Simple Texts for Text Simplification</i>	
Renliang Sun, Wei Xu and Xiaojun Wan	9345
<i>Acquiring Frame Element Knowledge with Deep Metric Learning for Semantic Frame Induction</i>	
Kosuke Yamada, Ryohei Sasano and Koichi Takeda	9356
<i>Leveraging Synthetic Targets for Machine Translation</i>	
Sarthak Mittal, Oleksii Hrinchuk and Oleksii Kuchaiev	9365
<i>Recipes for Sequential Pre-training of Multilingual Encoder and Seq2Seq Models</i>	
Saleh Soltan, Andy Rosenbaum, Tobias Falke, Qin Lu, Anna Rumshisky and Wael Hamza .	9380
<i>Constructing Code-mixed Universal Dependency Forest for Unbiased Cross-lingual Relation Extraction</i>	
Hao Fei, Meishan Zhang, Min Zhang and Tat-Seng Chua	9395
<i>Spontaneous gestures encoded by hand positions improve language models: An Information-Theoretic motivated study</i>	
Yang Xu and Yang Cheng	9409
<i>Progressive Translation: Improving Domain Robustness of Neural Machine Translation with Intermediate Sequences</i>	
Chaojun Wang, Yang Liu and Wai Lam	9425
<i>Controlled Text Generation with Hidden Representation Transformations</i>	
Vaibhav Kumar, Hana Koorehdavoudi, Masud Moshtaghi, Amita Misra, Ankit R Chadha and Emilio Ferrara	9440
<i>Visual Coherence Loss for Coherent and Visually Grounded Story Generation</i>	
Xudong Hong, Vera Demberg, Asad Sayeed, Qiankun Zheng and Bernt Schiele	9456
<i>AnaMeta: A Table Understanding Dataset of Field Metadata Knowledge Shared by Multi-dimensional Data Analysis Tasks</i>	
Xinyi He, Mengyu Zhou, Mingjie Zhou, Jialiang Xu, Xiao Lv, Tianle Li, Yijia Shao, Shi Han, Zejian Yuan and Dongmei Zhang	9471
<i>Large Language Models Are Partially Primed in Pronoun Interpretation</i>	
Suet-Ying Lam, Qingcheng Zeng, Kexun Zhang, Chenyu You and Rob Voigt	9493
<i>Counterfactuals of Counterfactuals: a back-translation-inspired approach to analyse counterfactual editors</i>	
George Filandrianos, Edmund G Dervakos, Orfeas Menis Mastromichalakis, Chrysoula Zerva and Giorgos Stamou	9507
<i>A Pilot Study on Dialogue-Level Dependency Parsing for Chinese</i>	
Gongyao Jiang, Shuang Liu, Meishan Zhang and Min Zhang	9526
<i>On the Off-Target Problem of Zero-Shot Multilingual Neural Machine Translation</i>	
Liang Chen, Shuming Ma, Dongdong Zhang, Furu Wei and Baobao Chang	9542

<i>ORCA: A Challenging Benchmark for Arabic Language Understanding</i>	
AbdelRahim Elmadany, ElMoatez Billah Nagoudi and Muhammad Abdul-Mageed	9559
<i>Delving into the Openness of CLIP</i>	
Shuhuai Ren, Lei Li, Xuancheng Ren, Guangxiang Zhao and Xu Sun	9587
<i>From Adversarial Arms Race to Model-centric Evaluation: Motivating a Unified Automatic Robustness Evaluation Framework</i>	
Yangyi Chen, Hongcheng Gao, Ganqu CUI, Lifan Yuan, Dehan Kong, Hanlu Wu, Ning Shi, Bo Yuan, Longtao Huang, hui xue, Zhiyuan Liu, Maosong Sun and Heng Ji	9607
<i>An Empirical Study of Sentiment-Enhanced Pre-Training for Aspect-Based Sentiment Analysis</i>	
Yice Zhang, Yifan Yang, Bin Liang, Shiwei Chen, Bing Qin and Ruifeng Xu	9633
<i>NatCS: Eliciting Natural Customer Support Dialogues</i>	
James Gung, Emily Moeng, Wesley Rose, Arshit Gupta, Yi Zhang and Saab Mansour	9652
<i>Are Intermediate Layers and Labels Really Necessary? A General Language Model Distillation Method</i>	
Shicheng Tan, Weng Lam Tam, Yuanchun Wang, Wenwen Gong, Shu Zhao, Peng Zhang and Jie Tang	9678
<i>Diable: Efficient Dialogue State Tracking as Operations on Tables</i>	
Pietro Lesci, Yoshinari Fujinuma, Momchil Hardalov, Chao Shang and Lluís Marquez	9697
<i>Neural Topic Modeling based on Cycle Adversarial Training and Contrastive Learning</i>	
Boyue Wang, Linhai Zhang, Deyu Zhou, Yi Cao and Jiandong Ding	9720
<i>Alleviating Exposure Bias via Multi-level Contrastive Learning and Deviation Simulation in Abstractive Summarization</i>	
Jiawen Xie, Qi Su, Shaoting Zhang and Xiaofan Zhang	9732
<i>Mapping Brains with Language Models: A Survey</i>	
Antonia Karamolegkou, Mostafa Abdou and Anders Søgaard	9748
<i>Parameter-Efficient Finetuning for Robust Continual Multilingual Learning</i>	
Kartikeya Badola, Shachi Dave and Partha Talukdar	9763
<i>Interpretable Multimodal Misinformation Detection with Logic Reasoning</i>	
Hui Liu, Wenya Wang and Haoliang Li	9781
<i>Semantic-conditioned Dual Adaptation for Cross-domain Query-based Visual Segmentation</i>	
Ye Wang, Tao Jin, Wang Lin, Xize Cheng, Linjun Li and Zhou Zhao	9797
<i>Figurative Language Processing: A Linguistically Informed Feature Analysis of the Behavior of Language Models and Humans</i>	
Hyewon Jang, Qi Yu and Diego Frassinelli	9816
<i>Taxonomy of Problems in Lexical Semantics</i>	
Bradley M Hauer and Grzegorz Kondrak	9833
<i>Making Pre-trained Language Models both Task-solvers and Self-calibrators</i>	
Yangyi Chen, Xingyao Wang and Heng Ji	9845
<i>EmbedTextNet: Dimension Reduction with Weighted Reconstruction and Correlation Losses for Efficient Text Embedding</i>	
Dae Yon Hwang, Bilal Taha and Yaroslav Nechaev	9863

<i>Denoising Enhanced Distantly Supervised Ultrafine Entity Typing</i> Yue Zhang, Hongliang Fei and Ping Li	9880
<i>INTapt: Information-Theoretic Adversarial Prompt Tuning for Enhanced Non-Native Speech Recognition</i> Eunseop Yoon, Hee Suk Yoon, John Harvill, Mark Hasegawa-Johnson and Chang D. Yoo ..	9893
<i>Local Temperature Beam Search: Avoid Neural Text DeGeneration via Enhanced Calibration</i> Dongkyu Lee, Gyeonghun KIM, Janghoon Han, Taesuk Hong, Yi-Reun Kim, Stanley Jungkyu Choi and Nevin L. Zhang	9903
<i>Explanation Graph Generation via Generative Pre-training over Synthetic Graphs</i> Han Cui, Shangzhan Li, Yu Zhang and Qi Shi	9916
<i>NaSGEC: a Multi-Domain Chinese Grammatical Error Correction Dataset from Native Speaker Texts</i> Yue Zhang, Bo Zhang, Haochen Jiang, Zhenghua Li, Chen Li, Fei Huang and Min Zhang ..	9935
<i>FORK: A Bite-Sized Test Set for Probing Culinary Cultural Biases in Commonsense Reasoning Models</i> Shramay Palta and Rachel Rudinger	9952
<i>FedPETuning: When Federated Learning Meets the Parameter-Efficient Tuning Methods of Pre-trained Language Models</i> Zhuo Zhang, Yuanhang Yang, Yong Dai, Qifan Wang, Yue Yu, Lizhen Qu and Zenglin Xu ..	9963
<i>MixPAVE: Mix-Prompt Tuning for Few-shot Product Attribute Value Extraction</i> Li Yang, Qifan Wang, Jingang Wang, Xiaojun Quan, Fuli Feng, Yu Chen, Madian Khabsa, Sinong Wang, Zenglin Xu and Dongfang Liu	9978
<i>SlowBERT: Slow-down Attacks on Input-adaptive Multi-exit BERT</i> Shengyao Zhang, Xudong Pan, Mi Zhang and Min Yang	9992
<i>Compositional Mathematical Encoding for Math Word Problems</i> Zhenwen Liang, Jipeng Zhang, Kehan Guo, Xiaodong Wu, Jie Shao and Xiangliang Zhang	10008
<i>PREADD: Prefix-Adaptive Decoding for Controlled Text Generation</i> Jonathan Pei, Kevin Yang and Dan Klein	10018
<i>EventOA: An Event Ontology Alignment Benchmark Based on FrameNet and Wikidata</i> Shaoru Guo, Chenhao Wang, Yubo Chen, Kang Liu, Ru Li and Jun Zhao	10038
<i>Enhancing Continual Relation Extraction via Classifier Decomposition</i> Heming Xia, Peiyi Wang, Tianyu Liu, Binghuai Lin, Yunbo Cao and Zhifang Sui	10053
<i>A Comparative Analysis of the Effectiveness of Rare Tokens on Creative Expression using ramBERT</i> Youbin Lee, Deokgi Kim, Byung-Won On and Ingyu Lee	10063
<i>MTR: A Dataset Fusing Inductive, Deductive, and Defeasible Reasoning</i> Yitian Li, Jidong Tian, Caoyun Fan, Wenqing Chen, Hao He and Yaohui Jin	10078
<i>NewsMet : A 'do it all' Dataset of Contemporary Metaphors in News Headlines</i> Rohan Joseph, Timothy Liu, Aik Beng Ng, Simon See and Sunny Rai	10090
<i>Concept2Box: Joint Geometric Embeddings for Learning Two-View Knowledge Graphs</i> Zijie Huang, Daheng Wang, Binxuan Huang, Chenwei Zhang, Jingbo Shang, Yan Liang, Zhen-gang Wang, Xian Li, Christos Faloutsos, Yizhou Sun and Wei Wang	10105

<i>Noise-Robust Training with Dynamic Loss and Contrastive Learning for Distantly-Supervised Named Entity Recognition</i>	
Zhiyuan Ma, Jintao Du and Shuheng Zhou	10119
<i>Take a Break in the Middle: Investigating Subgoals towards Hierarchical Script Generation</i>	
XINZE LI, Yixin Cao, Muhao Chen and Aixin Sun	10129
<i>End-to-End Task-Oriented Dialogue Systems Based on Schema</i>	
Wiradee Imrattana-trai and Ken Fukuda	10148
<i>HaVQA: A Dataset for Visual Question Answering and Multimodal Research in Hausa Language</i>	
Shantipriya Parida, Idris Abdulummin, Shamsuddeen Hassan Muhammad, Aneesh Bose, Guneet Singh Kohli, Ibrahim Said Ahmad, Ketan Kotwal, Sayan Deb Sarkar, Ondřej Bojar and Habeebah Adamu Kakudi	10162
<i>Claim-Dissector: An Interpretable Fact-Checking System with Joint Re-ranking and Veracity Prediction</i>	
Martin Fajcik, Petr Motliceck and Pavel Smrz	10184
<i>StructSP: Efficient Fine-tuning of Task-Oriented Dialog System by Using Structure-aware Boosting and Grammar Constraints</i>	
Truong Dinh Do, Phuong Minh Nguyen and Minh Le Nguyen	10206
<i>GDA: Generative Data Augmentation Techniques for Relation Extraction Tasks</i>	
Xuming Hu, Aiwei Liu, Zeqi Tan, Xin Zhang, Chenwei Zhang, Irwin King and Philip S. Yu	10221
<i>WebDP: Understanding Discourse Structures in Semi-Structured Web Documents</i>	
Peilin Liu, Hongyu Lin, Meng Liao, Hao Xiang, Xianpei Han and Le Sun	10235
<i>Tab-CoT: Zero-shot Tabular Chain of Thought</i>	
Jin Ziqi and Wei Lu	10259
<i>KNSE: A Knowledge-aware Natural Language Inference Framework for Dialogue Symptom Status Recognition</i>	
Wei Chen, Shiqi Wei, Zhongyu Wei and Xuanjing Huang	10278
<i>Augmenting Large Language Model Translators via Translation Memories</i>	
Yongyu Mu, Abudurexiti Reheman, Zhiquan Cao, Yuchun Fan, Bei Li, Yinqiao Li, Tong Xiao, Chunliang Zhang and Jingbo Zhu	10287
<i>Character Coreference Resolution in Movie Screenplays</i>	
Sabyasachee Baruah and Shrikanth Narayanan	10300
<i>Enhancing Event Causality Identification with Event Causal Label and Event Pair Interaction Graph</i>	
Ruili Pu, Yang Li, Suge Wang, Deyu Li, Jianxing Zheng and Jian Liao	10314
<i>LightFormer: Light-weight Transformer Using SVD-based Weight Transfer and Parameter Sharing</i>	
Xiuqing Lv, Peng Zhang, sunzhu Initial li, Guobing Gan and Yueheng Sun	10323
<i>Multi-hop Evidence Retrieval for Cross-document Relation Extraction</i>	
Keming Lu, I-Hung Hsu, Wenxuan Zhou, Mingyu Derek Ma and Muhao Chen	10336
<i>Which Examples Should be Multiply Annotated? Active Learning When Annotators May Disagree</i>	
Connor T Baumler, Anna Sotnikova and Hal Daumé III	10352
<i>PIP: Parse-Instructed Prefix for Syntactically Controlled Paraphrase Generation</i>	
Yixin Wan, Kuan-Hao Huang and Kai-Wei Chang	10372

<i>DePlot: One-shot visual language reasoning by plot-to-table translation</i>	
Fangyu Liu, Julian Martin Eisenschlos, Francesco Piccinno, Syrine Krichene, Chenxi Pang, Kenton Lee, Mandar Joshi, Wenhua Chen, Nigel Collier and Yasemin Altun	10381
<i>Stochastic Bridges as Effective Regularizers for Parameter-Efficient Tuning</i>	
Weize Chen, Xu Han, Yankai Lin, Zhiyuan Liu, Maosong Sun and Jie Zhou	10400
<i>Learning from a Friend: Improving Event Extraction via Self-Training with Feedback from Abstract Meaning Representation</i>	
Zhiyang Xu, Jay Yoon Lee and Lifu Huang	10421
<i>How Well Do Large Language Models Perform on Faux Pas Tests?</i>	
Natalie Shapira, Guy Zwirn and Yoav Goldberg	10438
<i>Modular Transformers: Compressing Transformers into Modularized Layers for Flexible Efficient Inference</i>	
Wangchunshu Zhou, Ronan Le Bras and Yejin Choi	10452
<i>ISLTranslate: Dataset for Translating Indian Sign Language</i>	
Abhinav Joshi, Susmit Agrawal and Ashutosh Modi	10466
<i>LMentry: A Language Model Benchmark of Elementary Language Tasks</i>	
Avia Efrat, Or Honovich and Omer Levy	10476
<i>Differentiable Instruction Optimization for Cross-Task Generalization</i>	
Masaru Isonuma, Junichiro Mori and Ichiro Sakata	10502
<i>Leveraging Training Data in Few-Shot Prompting for Numerical Reasoning</i>	
Zhanming Jie and Wei Lu	10518
<i>How does the task complexity of masked pretraining objectives affect downstream performance?</i>	
Atsuki Yamaguchi, Hiroaki Ozaki, Terufumi Morishita, Gaku Morio and Yasuhiro Sogawa	10527
<i>AUGUST: an Automatic Generation Understudy for Synthesizing Conversational Recommendation Datasets</i>	
Yu Lu, Junwei Bao, Zichen Ma, Xiaoguang Han, Youzheng Wu, Shuguang Cui and Xiaodong He	10538
<i>Knowing-how & Knowing-that: A New Task for Machine Comprehension of User Manuals</i>	
Hongru Liang, Jia Liu, Weihong Du, Dingnan Jin, Wenqiang Lei, Zujie Wen and Jiancheng Lv	10550
<i>Deep Span Representations for Named Entity Recognition</i>	
Enwei Zhu, Yiyang Liu and Jinpeng Li	10565
<i>Disambiguated Lexically Constrained Neural Machine Translation</i>	
Jinpeng Zhang, Nini Xiao, Ke Wang, Chuanqi Dong, Xiangyu Duan, Yuqi Zhang and Min Zhang	10583
<i>Curating Datasets for Better Performance with Example Training Dynamics</i>	
Aviad Sar-Shalom and Roy Schwartz	10597
<i>Multi-armed bandits for resource efficient, online optimization of language model pre-training: the use case of dynamic masking</i>	
Inigo Urteaga, Moulay Zaidane Draidia, Tomer Lancewicki and Shahram Khadivi	10609
<i>ERNIE-Code: Beyond English-Centric Cross-lingual Pretraining for Programming Languages</i>	
Yekun Chai, Shuohuan Wang, Chao Pang, Yu Sun, Hao Tian and Hua Wu	10628

<i>PromptAttack: Probing Dialogue State Trackers with Adversarial Prompts</i> Xiangjue Dong, Yun He, Ziwei Zhu and James Caverlee	10651
<i>Understanding Programs by Exploiting (Fuzzing) Test Cases</i> Jianyu Zhao, Yuyang Rong, Yiwen Guo, Yifeng He and Hao Chen	10667
<i>Hybrid Hierarchical Retrieval for Open-Domain Question Answering</i> Manoj Ghuhana Arivazhagan, Lan Liu, Peng Qi, Xinchu Chen, William Yang Wang and zhiheng huang	10680
<i>Coherent or Not? Stressing a Neural Language Model for Discourse Coherence in Multiple Languages</i> Dominique Brunato, Felice Dell’Orletta, irene dini and Andrea Amelio Ravelli	10690
<i>Understanding Differential Search Index for Text Retrieval</i> Xiaoyang Chen, Yanjiang Liu, Ben He, Le Sun and Yingfei Sun	10701
<i>Masked Audio Text Encoders are Effective Multi-Modal Rescorers</i> Jinglun Cai, Monica Sunkara, Xilai Li, Anshu Bhatia, Xiao Pan and Sravan Bodapati	10718
<i>Replace and Report: NLP Assisted Radiology Report Generation</i> Kaveri Kale, Pushpak Bhattacharyya and Kshitij Jadhav	10731
<i>Pre-trained Personalized Review Summarization with Effective Saliency Estimation</i> Hongyan Xu, Hongtao Liu, Zhepeng Lv, Qing Yang and Wenjun Wang	10743
<i>CaPE: Contrastive Parameter Ensembling for Reducing Hallucination in Abstractive Summarization</i> Prafulla Kumar Choubey, Alex Fabbri, Jesse Vig, Chien-Sheng Jason Wu, Wenhao Liu and Nazneen Rajani	10755
<i>OpineSum: Entailment-based self-training for abstractive opinion summarization</i> Annie Louis and Joshua Maynez	10774
<i>A Call for Standardization and Validation of Text Style Transfer Evaluation</i> Phil Sidney Ostheimer, Mayank Kumar Nagda, Marius Kloft and Sophie Fellenz	10791
<i>Bridging the Granularity Gap for Acoustic Modeling</i> Chen Xu, Yuhao Zhang, Chengbo Jiao, Xiaoqian Liu, Chi Hu, Xin Zeng, Tong Xiao, Anxiang Ma, Huizhen Wang and Jingbo Zhu	10816
<i>MMSD2.0: Towards a Reliable Multi-modal Sarcasm Detection System</i> Libo Qin, Shijue Huang, Qiguang Chen, Chenran Cai, Yudi Zhang, Bin Liang, Wanxiang Che and Ruifeng Xu	10834
<i>Learn to Not Link: Exploring NIL Prediction in Entity Linking</i> Fangwei Zhu, Jifan Yu, Hailong Jin, Lei Hou, Juanzi Li and Zhifang Sui	10846
<i>On Text-based Personality Computing: Challenges and Future Directions</i> Qixiang Fang, Anastasia Giachanou, Ayoub Bagheri, Laura Boeschoten, Erik-Jan van Kesteren, Mahdi Shafiee Kamalabad and Daniel Oberski	10861
<i>Structured Pruning for Efficient Generative Pre-trained Language Models</i> Chaofan Tao, Lu Hou, Haoli Bai, Jiansheng Wei, Xin Jiang, Qun Liu, Ping Luo and Ngai Wong	10880
<i>Prompt-Guided Retrieval Augmentation for Non-Knowledge-Intensive Tasks</i> Zhicheng Guo, Sijie Cheng, Yile Wang, Peng Li and Yang Liu	10896

<i>Contextualized Semantic Distance between Highly Overlapped Texts</i>	
Letian Peng, Zuchao Li and Hai Zhao	10913
<i>Unsupervised Dense Retrieval with Relevance-Aware Contrastive Pre-Training</i>	
Yibin Lei, Liang Ding, Yu Cao, Changtong Zan, Andrew Yates and Dacheng Tao	10932
<i>Verifying Annotation Agreement without Multiple Experts: A Case Study with Gujarati SNACS</i>	
Maitrey Mehta and Vivek Srikumar	10941
<i>Reinforced Active Learning for Low-Resource, Domain-Specific, Multi-Label Text Classification</i>	
Lukas Wertz, Jasmina Bogojeska, Katsiaryna Mirylenka and Jonas Kuhn	10959
<i>Improving Classroom Dialogue Act Recognition from Limited Labeled Data with Self-Supervised Contrastive Learning Classifiers</i>	
Vikram Kumaran, Jonathan Rowe, Bradford W Mott, Snigdha Chaturvedi and James Lester	10978
<i>Contrastive Token-Wise Meta-Learning for Unseen Performer Visual Temporal-Aligned Translation</i>	
Linjun Li, Tao Jin, Xize Cheng, Ye Wang, Wang Lin, Rongjie Huang and Zhou Zhao	10993
<i>Enhancing Cross-lingual Prompting with Dual Prompt Augmentation</i>	
Meng Zhou, Xin Li, Yue Jiang and Lidong Bing	11008
<i>Foveate, Attribute, and Rationalize: Towards Physically Safe and Trustworthy AI</i>	
Alex Mei, Sharon Levy and William Yang Wang	11021
<i>Multijugate Dual Learning for Low-Resource Task-Oriented Dialogue System</i>	
Shimin Li, Xiaotian Zhang, Yanjun Zheng, Linyang Li and Xipeng Qiu	11037
<i>A Class-Rebalancing Self-Training Framework for Distantly-Supervised Named Entity Recognition</i>	
Qi Li, Tingyu Xie, Peng Peng, Hongwei Wang and Gaoang Wang	11054
<i>MURMUR: Modular Multi-Step Reasoning for Semi-Structured Data-to-Text Generation</i>	
Swarnadeep Saha, Xinyan Velocity Yu, Mohit Bansal, Ramakanth Pasunuru and Asli Celikyilmaz	11069
<i>Learning by Analogy: Diverse Questions Generation in Math Word Problem</i>	
Zihao Zhou, Maizhen Ning, Qifeng Wang, Jie Yao, Wei Wang, xiaowei huang and Kaizhu Huang	11091
<i>Revisit Few-shot Intent Classification with PLMs: Direct Fine-tuning vs. Continual Pre-training</i>	
Haode Zhang, Haowen Liang, Li-Ming Zhan, Xiao-Ming Wu and Albert Y.S. Lam	11105
<i>Improving Contrastive Learning of Sentence Embeddings from AI Feedback</i>	
Qinyuan Cheng, Xiaogui Yang, Tianxiang Sun, Linyang Li and Xipeng Qiu	11122
<i>Mars: Modeling Context & State Representations with Contrastive Learning for End-to-End Task-Oriented Dialog</i>	
Haipeng Sun, Junwei Bao, Youzheng Wu and Xiaodong He	11139
<i>Text Augmented Open Knowledge Graph Completion via Pre-Trained Language Models</i>	
Pengcheng Jiang, Shivam Agarwal, Bowen Jin, Xuan Wang, Jimeng Sun and Jiawei Han ..	11161
<i>Discourse Analysis via Questions and Answers: Parsing Dependency Structures of Questions Under Discussion</i>	
Wei-Jen Ko, Yating Wu, Cutter J Dalton, Dananjay T Srinivas, Greg Durrett and Junyi Jessy Li	11181

<i>An Integrated Approach for Political Bias Prediction and Explanation Based on Discursive Structure</i> Nicolas Devatine, Philippe Muller and Chloé Braud	11196
<i>Smart Word Suggestions for Writing Assistance</i> Chenshuo Wang, Shaoguang Mao, Tao Ge, Wenshan Wu, Xun Wang, Yan Xia, Jonathan Tien and Dongyan Zhao	11212
<i>JECC: Commonsense Reasoning Tasks Derived from Interactive Fictions</i> Mo Yu, Yi Gu, Xiaoxiao Guo, Yufei Feng, Xiaodan Zhu, Michael Greenspan, Murray Campbell and Chuang Gan	11226
<i>A Study on Knowledge Distillation from Weak Teacher for Scaling Up Pre-trained Language Models</i> Hayeon Lee, Rui Hou, Jongpil Kim, Davis Liang, Sung Ju Hwang and Alexander Min	11239
<i>SORTIE: Dependency-Aware Symbolic Reasoning for Logical Data-to-text Generation</i> Xueliang Zhao, Tingchen Fu, Lemao Liu, Lingpeng Kong, Shuming Shi and Rui Yan	11247
<i>Boosting Event Extraction with Denoised Structure-to-Text Augmentation</i> bo wang, Heyan Huang, Xiaochi Wei, Ge Shi, Xiao Liu, Chong Feng, Tong Zhou, Shuaiqiang WANG and Dawei Yin	11267
<i>Detecting Adversarial Samples through Sharpness of Loss Landscape</i> Rui Zheng, Shihan Dou, Yuhao Zhou, Qin Liu, Tao Gui, Qi Zhang, Zhongyu Wei, Xuanjing Huang and Menghan Zhang	11282
<i>A Simple, Yet Effective Approach to Finding Biases in Code Generation</i> Spyridon Mouselinos, Mateusz Malinowski and Henryk Michalewski	11299
<i>Membership Inference Attacks against Language Models via Neighbourhood Comparison</i> Justus Mattern, Fatemehsadat Mireshghallah, Zhijing Jin, Bernhard Schoelkopf, Mrinmaya Sachan and Taylor Berg-Kirkpatrick	11330
<i>CFL: Causally Fair Language Models Through Token-level Attribute Controlled Generation</i> Rahul Madhavan, Rishabh Garg, Kahini Wadhawan and Sameep Mehta	11344
<i>Can Diffusion Model Achieve Better Performance in Text Generation ? Bridging the Gap between Training and Inference !</i> Zecheng Tang, Pinzheng Wang, Keyan Zhou, Juntao Li, Ziqiang Cao and Min Zhang	11359
<i>Topic-Guided Self-Introduction Generation for Social Media Users</i> Chunpu Xu, Jing Li, Piji Li and Min Yang	11387
<i>Recyclable Tuning for Continual Pre-training</i> Yujia Qin, Cheng Qian, Xu Han, Yankai Lin, Huadong Wang, Ruobing Xie, Zhiyuan Liu, Mao-song Sun and Jie Zhou	11403
<i>BLOCSUM: Block Scope-based Source Code Summarization via Shared Block Representation</i> YunSeok Choi, Hyojun Kim and Jee-Hyong Lee	11427
<i>HyperPELT: Unified Parameter-Efficient Language Model Tuning for Both Language and Vision-and-Language Tasks</i> Zhengkun Zhang, Wenya Guo, Xiaojun Meng, Yasheng Wang, Yadao Wang, Xin Jiang, Qun Liu and Zhenglu Yang	11442
<i>Enhancing Unsupervised Semantic Parsing with Distributed Contextual Representations</i> Zixuan Ling, Xiaoqing Zheng, Jianhan Xu, Jinshu Lin, Kai-Wei Chang, Cho-Jui Hsieh and Xuanjing Huang	11454

<i>Generating Labeled Data for Relation Extraction: A Meta Learning Approach with Joint GPT-2 Training</i>	
Amir Pouran Ben Veysseh, Franck Dernoncourt, Bonan Min and Thien Huu Nguyen	11466
<i>Disfluency Generation for More Robust Dialogue Systems</i>	
Benjamin Marie	11479
<i>Dipping PLMs Sauce: Bridging Structure and Text for Effective Knowledge Graph Completion via Conditional Soft Prompting</i>	
Chen Chen, Yufei Wang, Aixin Sun, Bing Li and Kwok-Yan Lam	11489
<i>Revisiting Pathologies of Neural Models under Input Reduction</i>	
Canasai Kruengkrai and Junichi Yamagishi	11504
<i>Lego-MT: Learning Detachable Models for Massively Multilingual Machine Translation</i>	
Fei Yuan, Yinquan Lu, Wenhao Zhu, Lingpeng Kong, Lei Li, Yu Qiao and Jingjing Xu . . .	11518
<i>FiDO: Fusion-in-Decoder optimized for stronger performance and faster inference</i>	
Michiel de Jong, Yury Zemlyanskiy, Joshua Ainslie, Nicholas FitzGerald, Sumit Sanghai, Fei Sha and William Cohen	11534
<i>Detecting Edit Failures In Large Language Models: An Improved Specificity Benchmark</i>	
Jason Hoelscher-Obermaier, Julia H Persson, Esben Kran, Ioannis Konostas and Fazl Barez	11548
<i>Structure-Aware Language Model Pretraining Improves Dense Retrieval on Structured Data</i>	
Xinze Li, Zhenghao Liu, Chenyan Xiong, Shi Yu, Yu Gu, Zhiyuan Liu and Ge Yu	11560
<i>Few-shot Joint Multimodal Aspect-Sentiment Analysis Based on Generative Multimodal Prompt</i>	
Xiaocui Yang, Shi Feng, Daling Wang, Qi Sun, Wenfang Wu, Yifei Zhang, Pengfei Hong and Soujanya Poria	11575
<i>Predicting Human Translation Difficulty Using Automatic Word Alignment</i>	
Zheng Wei Lim, Trevor Cohn, Charles Kemp and Ekaterina Vylomova	11590
<i>Know Where You're Going: Meta-Learning for Parameter-Efficient Fine-Tuning</i>	
Mozhdeh Gheini, Xuezhe Ma and Jonathan May	11602
<i>Moving Beyond Downstream Task Accuracy for Information Retrieval Benchmarking</i>	
Keshav Santhanam, Jon Saad-Falcon, Martin Franz, Omar Khattab, Avi Sil, Radu Florian, Md Arafat Sultan, Salim Roukos, Matei Zaharia and Christopher Potts	11613
<i>AxomiyaBERTa: A Phonologically-aware Transformer Model for Assamese</i>	
Abhijnan Nath, Sheikh Abdul Mannan and Nikhil Krishnaswamy	11629
<i>An Exploratory Study on Model Compression for Text-to-SQL</i>	
Shuo Sun, Yuze Gao, Yuchen Zhang, Jian Su, Bin Chen, Yingzhan Lin and Shuqi Sun . . .	11647
<i>FluentSpeech: Stutter-Oriented Automatic Speech Editing with Context-Aware Diffusion Models</i>	
Ziyue Jiang, Qian Yang, Jialong Zuo, Zhenhui Ye, Rongjie Huang, Yi Ren and Zhou Zhao	11655
<i>HyHTM: Hyperbolic Geometry-based Hierarchical Topic Model</i>	
Simra Shahid, Tanay Anand, Nikitha Srikanth, Sumit Bhatia, Balaji Krishnamurthy and nikaash puri	11672
<i>KoRC: Knowledge Oriented Reading Comprehension Benchmark for Deep Text Understanding</i>	
Zijun Yao, Yantao Liu, Xin Lv, Shulin Cao, Jifan Yu, Juanzi Li and Lei Hou	11689

<i>DKAF: KB Arbitration for Learning Task-Oriented Dialog Systems with Dialog-KB Inconsistencies</i> Vishal Vivek Saley, Rocktim Jyoti Das, Dinesh Raghu and Mausam -	11708
<i>Scale-Invariant Infinite Hierarchical Topic Model</i> Shusei Eshima and Daichi Mochihashi	11731
<i>RC3: Regularized Contrastive Cross-lingual Cross-modal Pre-training</i> Chulun Zhou, Yunlong Liang, Fandong Meng, Jinan Xu, Jinsong Su and Jie Zhou	11747
<i>Deep Equilibrium Non-Autoregressive Sequence Learning</i> Zaixiang Zheng, Yi Zhou and Hao Zhou	11763
<i>ReGen: Zero-Shot Text Classification via Training Data Generation with Progressive Dense Retrieval</i> Yue Yu, Yuchen Zhuang, Rongzhi Zhang, Yu Meng, Jiaming Shen and Chao Zhang	11782
<i>Race, Gender, and Age Biases in Biomedical Masked Language Models</i> Michelle YoungJin Kim, Junghwan Kim and Kristen Johnson	11806
<i>Neighboring Words Affect Human Interpretation of Saliency Explanations</i> Alon Jacovi, Hendrik Schuff, Heike Adel, Ngoc Thang Vu and Yoav Goldberg	11816
<i>HELP ME THINK: A Simple Prompting Strategy for Non-experts to Create Customized Content with Models</i> Swaroop Mishra and Elnaz Nouri	11834
<i>Decker: Double Check with Heterogeneous Knowledge for Commonsense Fact Verification</i> Anni Zou, Zhuosheng Zhang and Hai Zhao	11891
<i>DopplerBAS: Binaural Audio Synthesis Addressing Doppler Effect</i> Jinglin Liu, Zhenhui Ye, Qian Chen, Siqi Zheng, Wen Wang, Zhang Qinglin and Zhou Zhao	11905
<i>Easy-to-Hard Learning for Information Extraction</i> Chang Gao, Wenxuan Zhang, Wai Lam and Lidong Bing	11913
<i>SConE: Simplified Cone Embeddings with Symbolic Operators for Complex Logical Queries</i> Chau Duc Minh Nguyen, Tim N French, Wei Liu and Michael Stewart	11931
<i>Two Heads Are Better Than One: Improving Fake News Video Detection by Correlating with Neighbors</i> Peng Qi, Yuyang Zhao, Yufeng Shen, Wei Ji, Juan Cao and Tat-Seng Chua	11947
<i>An Annotated Dataset for Explainable Interpersonal Risk Factors of Mental Disturbance in Social Media Posts</i> Muskan Garg, Amirmohammad Shahbandegan, Amrit Chadha and Vijay Mago	11960
<i>Nano: Nested Human-in-the-Loop Reward Learning for Few-shot Language Model Control</i> Xiang Fan, Yiwei Lyu, Paul Pu Liang, Ruslan Salakhutdinov and Louis-Philippe Morency	11970
<i>Connectivity Patterns are Task Embeddings</i> Zhiheng Xi, Rui Zheng, Yuansen Zhang, Xuanjing Huang, Zhongyu Wei, Minlong Peng, Mingming Sun, Qi Zhang and Tao Gui	11993
<i>Improving Autoregressive Grammatical Error Correction with Non-autoregressive Models</i> Hang Cao, Zhiquan Cao, Chi Hu, Baoyu Hou, Tong Xiao and Jingbo Zhu	12014
<i>SamToNe: Improving Contrastive Loss for Dual Encoder Retrieval Models with Same Tower Negatives</i> Fedor Moiseev, Gustavo Hernandez Abrego, Peter Dornbach, Imed Zitouni, Enrique Alfonseca and Zhe Dong	12028

<i>On the Strength of Sequence Labeling and Generative Models for Aspect Sentiment Triplet Extraction</i> Shen Zhou and Tiejun Qian	12038
<i>Revisiting Non-Autoregressive Translation at Scale</i> Zhihao Wang, Longyue Wang, Jinsong Su, Junfeng Yao and Zhaopeng Tu	12051
<i>Improving Radiology Summarization with Radiograph and Anatomy Prompts</i> Jinpeng hu, Zhihong Chen, Yang Liu, Xiang Wan and Tsung-Hui Chang	12066
<i>Explanation Regeneration via Information Bottleneck</i> Qintong Li, Zhiyong Wu, Lingpeng Kong and Wei Bi	12081
<i>Improving Zero-shot Multilingual Neural Machine Translation by Leveraging Cross-lingual Consistency Regularization</i> Pengzhi Gao, Liwen Zhang, Zhongjun He, Hua Wu and Haifeng Wang	12103
<i>ReactIE: Enhancing Chemical Reaction Extraction with Weak Supervision</i> Ming Zhong, Siru Ouyang, Minhao Jiang, Vivian Hu, Yizhu Jiao, Xuan Wang and Jiawei Han	12120
<i>Expand, Rerank, and Retrieve: Query Reranking for Open-Domain Question Answering</i> Yung-Sung Chuang, Wei Fang, Shang-Wen Li, Wen-tau Yih and James Glass	12131
<i>Neural Networks Against (and For) Self-Training: Classification with Small Labeled and Large Unlabeled Sets</i> Payam Karisani	12148
<i>Inducing Character-level Structure in Subword-based Language Models with Type-level Interchange Intervention Training</i> Jing Huang, Zhengxuan Wu, Kyle Mahowald and Christopher Potts	12163
<i>Efficient Document Embeddings via Self-Contrastive Bregman Divergence Learning</i> Daniel Saggau, Mina Rezaei, Bernd Bischl and Ilias Chalkidis	12181
<i>QAP: A Quantum-Inspired Adaptive-Priority-Learning Model for Multimodal Emotion Recognition</i> Ziming Li, Yan Zhou, Yaxin Liu, Fuqing Zhu, Chuanpeng Yang and Songlin Hu	12191
<i>Language acquisition: do children and language models follow similar learning stages?</i> Linnea Evanson, Yair Lakretz and Jean Rémi king	12205
<i>The Role of Output Vocabulary in T2T LMs for SPARQL Semantic Parsing</i> Debayan Banerjee, Pranav Ajit Nair, Ricardo Usbeck and Chris Biemann	12219
<i>UniCOQE: Unified Comparative Opinion Quintuple Extraction As A Set</i> Zinong Yang, Feng Xu, Jianfei Yu and Rui Xia	12229
<i>Response-conditioned Turn-taking Prediction</i> Bing'er Jiang, Erik Ekstedt and Gabriel Skantze	12241
<i>A Unified One-Step Solution for Aspect Sentiment Quad Prediction</i> Junxian Zhou, Haiqin Yang, Yuxuan He, Hao Mou and JunBo Yang	12249
<i>On Isotropy, Contextualization and Learning Dynamics of Contrastive-based Sentence Representation Learning</i> Chenghao Xiao, Yang Long and Noura Al Moubayed	12266
<i>Few-shot Fine-tuning vs. In-context Learning: A Fair Comparison and Evaluation</i> Marius Mosbach, Tiago Pimentel, Shauli Ravfogel, Dietrich Klakow and Yanai Elazar	12284

<i>Common Law Annotations: Investigating the Stability of Dialog System Output Annotations</i>	
Seunggun Lee, Alexandra DeLucia, Nikita Nangia, Praneeth S Ganedi, Ryan Min Guan, Rubing Li, Britney A Ngaw, Aditya Singhal, Shalaka Vaidya, Zijun Yuan, Lining Zhang and João Sedoc	12315
<i>HuaSLIM: Human Attention Motivated Shortcut Learning Identification and Mitigation for Large Language models</i>	
Yuqi Ren and Deyi Xiong	12350
<i>PMI-Align: Word Alignment With Point-Wise Mutual Information Without Requiring Parallel Training Data</i>	
Fatemeh Azadi, Hesham Faili and Mohammad Javad Dousti	12366
<i>Exploring Non-Verbal Predicates in Semantic Role Labeling: Challenges and Opportunities</i>	
Riccardo Orlando, Simone Conia and Roberto Navigli	12378
<i>DSPM-NLG: A Dual Supervised Pre-trained Model for Few-shot Natural Language Generation in Task-oriented Dialogue System</i>	
Yufan Wang, Bowei Zou, Rui Fan, Ai Ti Aw and Tingting He	12389
<i>TEPrompt: Task Enlightenment Prompt Learning for Implicit Discourse Relation Recognition</i>	
Wei Xiang, Chao Liang and Bang Wang	12403
<i>Evaluating Factuality in Cross-lingual Summarization</i>	
Mingqi Gao, Wenqing Wang, Xiaojun Wan and Yuemei Xu	12415
<i>On the Correspondence between Compositionality and Imitation in Emergent Neural Communication</i>	
Emily Cheng, Mathieu Rita and Thierry Poibeau	12432
<i>The Coreference under Transformation Labeling Dataset: Entity Tracking in Procedural Texts Using Event Models</i>	
Kyeongmin Rim, Jingxuan Tu, Bingyang Ye, Marc Verhagen, Eben Holderness and James Pustejovsky	12448
<i>Why Does Zero-Shot Cross-Lingual Generation Fail? An Explanation and a Solution</i>	
Tianjian Li and Kenton Murray	12461
<i>Distractor Generation based on Text2Text Language Models with Pseudo Kullback-Leibler Divergence Regulation</i>	
Hui-Juan Wang, Kai-Yu Hsieh, Han-Cheng Yu, Jui-Ching Tsou, Yu An Shih, Chen-Hua Huang and Yao-Chung Fan	12477
<i>Lexical Translation Inconsistency-Aware Document-Level Translation Repair</i>	
Zhen Zhang, Junhui Li, Shimin Tao and Hao Yang	12492
<i>CausalDialogue: Modeling Utterance-level Causality in Conversations</i>	
Yi-Lin Tuan, Alon Albalak, Wenda Xu, Michael S Saxon, Connor F Pryor, Lise Getoor and William Yang Wang	12506
<i>Towards Unified Spoken Language Understanding Decoding via Label-aware Compact Linguistics Representations</i>	
Zhihong Zhu, Xuxin Cheng, Zhiqi Huang, Dongsheng Chen and Yuexian Zou	12523
<i>Less Likely Brainstorming: Using Language Models to Generate Alternative Hypotheses</i>	
Liyan Tang, Yifan Peng, Yanshan Wang, Ying Ding, Greg Durrett and Justin F Rousseau	12532
<i>Language Modeling with Latent Situations</i>	
Belinda Z. Li, Maxwell Nye and Jacob Andreas	12556

<i>Can Cross-Lingual Transferability of Multilingual Transformers Be Activated Without End-Task Data?</i> Zewen Chi, Heyan Huang and Xian-Ling Mao	12572
<i>Focus-aware Response Generation in Inquiry Conversation</i> Yiquan Wu, Weiming Lu, Yating Zhang, Adam Jatowt, Jun Feng, Changlong Sun, Fei Wu and Kun Kuang	12585
<i>A Hierarchical Explanation Generation Method Based on Feature Interaction Detection</i> Yiming Ju, Yuanzhe Zhang, Kang Liu and Jun Zhao	12600
<i>Jointly Reparametrized Multi-Layer Adaptation for Efficient and Private Tuning</i> Umang Gupta, Aram Galstyan and Greg Ver Steeg	12612
<i>A Diffusion Model for Event Skeleton Generation</i> Fangqi Zhu, Lin Zhang, Jun Gao, Bing Qin, Ruifeng Xu and Haiqin Yang	12630
<i>Nonparametric Decoding for Generative Retrieval</i> Hyunji Lee, JaeYoung Kim, Hoyeon Chang, hanseok Oh, Sohee Yang, Vladimir Karpukhin, Yi Lu and Minjoon Seo	12642
<i>Aspect-aware Unsupervised Extractive Opinion Summarization</i> Haoyuan Li, Somnath Basu Roy Chowdhury and Snigdha Chaturvedi	12662
<i>GNN-SL: Sequence Labeling Based on Nearest Examples via GNN</i> Shuhe Wang, Yuxian Meng, Rongbin Ouyang, Jiwei Li, Tianwei Zhang, Lingjuan Lyu and Guoyin Wang	12679
<i>Serial Contrastive Knowledge Distillation for Continual Few-shot Relation Extraction</i> Xinyi Wang, Zitao Wang and Wei Hu	12693
<i>Revisiting the Architectures like Pointer Networks to Efficiently Improve the Next Word Distribution, Summarization Factuality, and Beyond</i> Haw-Shiuan Chang, Zonghai Yao, Alolika Gon, hong yu and Andrew McCallum	12707
<i>GLUE-X: Evaluating Natural Language Understanding Models from an Out-of-Distribution Generali- zation Perspective</i> Linyi Yang, Shuibai Zhang, Libo Qin, Yafu Li, Yidong Wang, Hanmeng Liu, Jindong Wang, Xing Xie and Yue Zhang	12731
<i>Investigating the Saliency of Sentiment Expressions in Aspect-Based Sentiment Analysis</i> Joachim Wagner and Jennifer Foster	12751
<i>DMLM: Descriptive Masked Language Modeling</i> Edoardo Barba, Niccolò Campolungo and Roberto Navigli	12770
<i>Reproducibility in NLP: What Have We Learned from the Checklist?</i> Ian Magnusson, Noah A. Smith and Jesse Dodge	12789
<i>Domain Generalization via Switch Knowledge Distillation for Robust Review Representation</i> You Zhang, Jin Wang, Liang-Chih Yu, Dan Xu and Xuejie Zhang	12812
<i>On Search Strategies for Document-Level Neural Machine Translation</i> Christian Herold and Hermann Ney	12827
<i>Causal Intervention for Mitigating Name Bias in Machine Reading Comprehension</i> Jiazheng Zhu, shaojuan wu, Xiaowang Zhang, Yuexian Hou and Zhiyong Feng	12837

<i>Counterfactual Probing for the Influence of Affect and Specificity on Intergroup Bias</i> Venkata Subrahmanyam Govindarajan, David I Beaver, Kyle Mahowald and Junyi Jessy Li	12853
<i>SongRewriter: A Chinese Song Rewriting System with Controllable Content and Rhyme Scheme</i> Yusen Sun, Liangyou Li, Qun Liu and Dit-Yan Yeung	12863
<i>Triplet-Free Knowledge-Guided Response Generation</i> Dongming Li, Jianfeng Liu and Baoyuan Wang	12881
<i>Implicit Memory Transformer for Computationally Efficient Simultaneous Speech Translation</i> Matthew Raffel and Lizhong Chen	12900
<i>Enhancing Document-level Event Argument Extraction with Contextual Clues and Role Relevance</i> Wanlong Liu, shaohuan cheng, Dingyi Zeng and Qu Hong	12908
<i>Exploring the Impact of Vision Features in News Image Captioning</i> Junzhe Zhang and Xiaojun Wan	12923
<i>Using Collostructional Analysis to evaluate BERT's representation of linguistic constructions</i> Tim Veenboer and Jelke Bloem	12937
<i>Selecting Better Samples from Pre-trained LLMs: A Case Study on Question Generation</i> Xingdi Yuan, Tong Wang, Yen-Hsiang Wang, Emery Fine, Rania Abdelghani, Hélène Sauzéon and Pierre-Yves Oudeyer	12952
<i>Sentiment Knowledge Enhanced Self-supervised Learning for Multimodal Sentiment Analysis</i> Fan Qian, Jiqing Han, Yongjun He, Tieran Zheng and Guibin Zheng	12966
<i>Theory of Mind in Freely-Told Children's Narratives: A Classification Approach</i> Bram van Dijk, Marco Spruit and Max J. van Duijn	12979
<i>Better Language Models of Code through Self-Improvement</i> Hung Quoc To, Nghi D. Q. Bui, Jin L.C. Guo and Tien N Nguyen	12994
<i>Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them</i> Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed H Chi, Denny Zhou and Jason Wei	13003
<i>Score It All Together: A Multi-Task Learning Study on Automatic Scoring of Argumentative Essays</i> Yuning Ding, Marie Bexte and Andrea Horbach	13052
<i>Data Sampling and (In)stability in Machine Translation Evaluation</i> Chi-kiu Lo and Rebecca Knowles	13064
<i>Probing Graph Decomposition for Argument Pair Extraction</i> Yang Sun, Bin Liang, Jianzhu Bao, Yice Zhang, Geng Tu, Min Yang and Ruifeng Xu	13075
<i>DiffuSum: Generation Enhanced Extractive Summarization with Diffusion</i> Haopeng Zhang, Xiao Liu and Jiawei Zhang	13089
<i>Towards Parameter-Efficient Integration of Pre-Trained Language Models In Temporal Video Grounding</i> Erica Kido Shimomoto, Edison Marrese-Taylor, Hiroya Takamura, Ichiro Kobayashi, Hideki Nakayama and Yusuke Miyao	13101
<i>A Memory Model for Question Answering from Streaming Data Supported by Rehearsal and Anticipation of Coreference Information</i> Vladimir Araujo, Alvaro M Soto and Marie-Francine Moens	13124

<i>Pay Attention to Implicit Attribute Values: A Multi-modal Generative Framework for AVE Task</i> Yupeng Zhang, Shensi Wang, Peiguang Li, Guanting Dong, Sirui Wang, Yunsen Xian, Zhoujun Li and Hongzhi Zhang	13139
<i>CoRRPUS: Code-based Structured Prompting for Neurosymbolic Story Understanding</i> Yijiang Dong, Lara J. Martin and Chris Callison-Burch	13152
<i>Fighting Bias With Bias: Promoting Model Robustness by Amplifying Dataset Biases</i> Yuval Reif and Roy Schwartz	13169
<i>Context-Aware Document Simplification</i> Liam Cripwell, Joël Legrand and Claire Gardent	13190
<i>Distinguish Before Answer: Generating Contrastive Explanation as Knowledge for Commonsense Question Answering</i> Qianglong Chen, Guohai Xu, Ming Yan, Ji Zhang, Fei Huang, Luo Si and Yin Zhang	13207
<i>Abstract then Play: A Skill-centric Reinforcement Learning Framework for Text-based Games</i> Anjie Zhu, Peng-Fei Zhang, Yi Zhang, Zi Huang and Jie Shao	13225
<i>SSP: Self-Supervised Post-training for Conversational Search</i> Quan Tu, Shen Gao, Xiaolong Wu, Zhao Cao, Ji-Rong Wen and Rui Yan	13237
<i>Towards Reference-free Text Simplification Evaluation with a BERT Siamese Network Architecture</i> Xinran Zhao, Esin Durmus and Dit-Yan Yeung	13250
<i>Causal interventions expose implicit situation models for commonsense language understanding</i> Takateru Yamakoshi, James L McClelland, Adele Goldberg and Robert Hawkins	13265
<i>Iterative Nearest Neighbour Machine Translation for Unsupervised Domain Adaptation</i> Hui Huang, Shuangzhi Wu, Xinnian Liang, Zefan Zhou, Muyun Yang and Tiejun Zhao	13294
<i>PruMux: Augmenting Data Multiplexing with Model Compression</i> Yushan Su, Vishvak Murahari, Karthik Narasimhan and Kai Li	13302
<i>With Prejudice to None: A Few-Shot, Multilingual Transfer Learning Approach to Detect Social Bias in Low Resource Languages</i> Nihar Ranjan Sahoo, Niteesh Kumar Reddy Mallela and Pushpak Bhattacharyya	13316
<i>Don't Lose Yourself! Empathetic Response Generation via Explicit Self-Other Awareness</i> Weixiang Zhao, Yanyan Zhao, Xin Lu and Bing Qin	13331
<i>Are Layout-Infused Language Models Robust to Layout Distribution Shifts? A Case Study with Scientific Documents</i> Catherine Chen, Zejiang Shen, Dan Klein, Gabriel Stanovsky, Doug Downey and Kyle Lo	13345
<i>Enhancing Neural Topic Model with Multi-Level Supervisions from Seed Words</i> Yang Lin, Xin Gao, Xu Chu, Yasha Wang, Junfeng Zhao and Chao Chen	13361
<i>Learning from Children: Improving Image-Caption Pretraining via Curriculum</i> Hammad Ayyubi, Rahul Lokesh, Alireza Zareian, Bo Wu and Shih-Fu Chang	13378
<i>Discovering Language Model Behaviors with Model-Written Evaluations</i> Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig R Pettit, Catherine Olsson, Sandipan Kundu, Saurav K Kadavath, Andy Jones, Anna Chen, Benjamin Mann, Brian R Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela	

Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Guro Khundadze, Jackson Kernion, James Landis, Jamie D Kerr, Jared Mueller, Jeeyoon Hyun, Joshua Lee Landau, kamal ndousse, Landon Goldberg, Liane Lovitt, Martin Lucas, Michael Sellitto, Miranda Huixin Zhang, Neerav Kingsland, Nelson Elhage, Nicholas Joseph, Noemi Mercado, Nova Xenia DasSarma, Oliver Rausch, Robin Larson, Sam McCandlish, Scott Johnston, Shauna Kravec, Sheer El Showk, Tamera Lanham, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Glen Hume, Yuntao Bai, Zac Hatfield-Dodds, Jack Clark, Samuel R. Bowman, Amanda Askell, Roger Baker Grosse, Danny Hernandez, Deep Ganguli, Evan J Hubinger, Nicholas Schiefer and Jared Kaplan 13387

Cross-Domain Argument Quality Estimation
 Michael Fromm, Max Berrendorf, Evgeniy Faerman and Thomas Seidl 13435

DiaASQ: A Benchmark of Conversational Aspect-based Sentiment Quadruple Analysis
 Bobo Li, Hao Fei, Fei Li, Yuhan Wu, Jinsong Zhang, Shengqiong Wu, Jingye Li, Yijiang Liu, Lizi Liao, Tat-Seng Chua and Donghong Ji 13449

GeoDRL: A Self-Learning Framework for Geometry Problem Solving using Reinforcement Learning in Deductive Reasoning
 Shuai Peng, Di Fu, Yijun Liang, Liangcai Gao and Zhi TANG 13468

Uncertainty-Aware Unlikelihood Learning Improves Generative Aspect Sentiment Quad Prediction
 Mengting Hu, Yinhao Bai, Yike Wu, Zhen Zhang, Liqi Zhang, Hang Gao, Shiwan Zhao and Minlie Huang 13481

Adversarial Knowledge Stimulated Contrastive Prompting for Few-shot Language Learners
 Kai Zheng, Qingfeng Sun, Yaming Yang, Tengchao Lv, Yeyong Pi, changlin zhao, Fei Xu and Qi Zhang 13495

Making Pre-trained Language Models Better Learn Few-Shot Spoken Language Understanding in More Practical Scenarios
 Yufan Wang, Jie Mei, Bowei Zou, Rui Fan, Tingting He and Ai Ti Aw 13508

Typology Guided Multilingual Position Representations: Case on Dependency Parsing
 Tao Ji, Yuanbin Wu and xiaoling Wang 13524

Learning Event-aware Measures for Event Coreference Resolution
 Yao Yao, Zuchao Li and Hai Zhao 13542

Second Language Acquisition of Neural Language Models
 Miyu Oba, Tatsuki Kuribayashi, Hiroki Ouchi and Taro Watanabe 13557

On the Universal Adversarial Perturbations for Efficient Data-free Adversarial Detection
 SongYang Gao, Shihan Dou, Qi Zhang, Xuanjing Huang, Jin Ma and Ying Shan 13573

Exploring the Effectiveness of Prompt Engineering for Legal Reasoning Tasks
 Fangyi Yu, Lee Quartey and Frank Schilder 13582

End-to-end Aspect-based Sentiment Analysis with Combinatory Categorical Grammar
 Yuanhe Tian, Weidong Chen, Bo Hu, Yan Song and Fei Xia 13597

ConKI: Contrastive Knowledge Injection for Multimodal Sentiment Analysis
 Yakun Yu, Mingjun Zhao, Shi-ang Qi, Feiran Sun, Baoxun Wang, Weidong Guo, Xiaoli Wang, Lei Yang and Di Niu 13610

On Degrees of Freedom in Defining and Testing Natural Language Understanding
 Saku Sugawara and Shun Tsugita 13625

<i>AttenWalker: Unsupervised Long-Document Question Answering via Attention-based Graph Walking</i> Yuxiang Nie, Heyan Huang, Wei Wei and Xian-Ling Mao	13650
<i>Adaptive Ordered Information Extraction with Deep Reinforcement Learning</i> Wenhao Huang, Jiaqing Liang, Zhixu Li, Yanghua Xiao and chuanjun ji	13664
<i>Wasserstein-Fisher-Rao Embedding: Logical Query Embeddings with Local Comparison and Global Transport</i> Zihao Wang, Weizhi Fei, Hang Yin, Yangqiu Song, Ginny Y. Wong and Simon See	13679
<i>RISE: Leveraging Retrieval Techniques for Summarization Evaluation</i> David Uthus and Jianmo Ni	13697
<i>On the Difference of BERT-style and CLIP-style Text Encoders</i> Zhihong Chen, Guiming Hardy Chen, Shizhe Diao, Xiang Wan and Benyou Wang	13710
<i>Model Interpretability and Rationale Extraction by Input Mask Optimization</i> Marc Felix Brinner and Sina Zarrieß	13722
<i>NusaCrowd: Open Source Initiative for Indonesian NLP Resources</i> Samuel Cahyawijaya, Holy Lovenia, Alham Fikri Aji, Genta Indra Winata, Bryan Wilie, Fajri Koto, Rahmad Mahendra, Christian Wibisono, Ade Romadhony, Karissa Vincentio, Jennifer Santoso, David Moeljadi, Cahya Wirawan, Frederikus Hudi, Muhammad Satrio Wicaksono, Ivan Halim Parnomanan, Ika Alfina, Ilham Firdausi Putra, Samsul Rahmadani, Yulianti Oenang, Ali Akbar Septiandri, James Jaya, Kaustubh Dhole, Arie Suryani, Rifki Afina Putri, Dan Su, Keith David Stevens, Made Nindyatama Nityasya, Muhammad Farid Adilazuarda, Ryan Ignatius Hadiwijaya, Ryandito Diandaru, Tiezheng Yu, Vito Ghifari, Wenliang Dai, Yan Xu, Dyah Inastra Damapuspita, Haryo Akbarianto Wibowo, Cuk Tho, Ichwanul Muslim Karo Karo, Tirana Noor Fatyanosa, Ziwei Ji, Graham Neubig, Timothy Baldwin, Sebastian Ruder, Pascale Fung, Herry Sujaini, Sakriani Sakti and Ayu Purwarianti	13745
<i>Transcribing Vocal Communications of Domestic Shiba Inu Dogs</i> Jieyi Huang, Chunhao Zhang, Mengyue Wu and Kenny Q. Zhu	13819
<i>SkillQG: Learning to Generate Question for Reading Comprehension Assessment</i> Xiaoqiang Wang, Bang Liu, Siliang Tang and Lingfei Wu	13833
<i>Improving Long Dialogue Summarization with Semantic Graph Representation</i> Yilun Hua, Zhaoyuan Deng and Kathleen McKeown	13851
<i>Model Intrinsic Features of Fine-tuning based Text Summarization Models for Factual Consistency</i> Jongyoon Song, Nohil Park, Bongkyu Hwang, Jaewoong Yun, Seongho Joe, Youngjune L Gwon and Sungroh Yoon	13884
<i>EfficientVLM: Fast and Accurate Vision-Language Models via Knowledge Distillation and Modal-adaptive Pruning</i> Tiannan Wang, Wangchunshu Zhou, Yan Zeng and Xinsong Zhang	13899
<i>DP-BART for Privatized Text Rewriting under Local Differential Privacy</i> Timour Igamberdiev and Ivan Habernal	13914
<i>Robustness of Learning from Task Instructions</i> Jiasheng Gu, Hongyu Zhao, Hanzi Xu, Liangyu Nie, Hongyuan Mei and Wenpeng Yin ...	13935
<i>Masked Latent Semantic Modeling: an Efficient Pre-training Alternative to Masked Language Modeling</i> Gábor Berend	13949

<i>Detection and Mitigation of the Negative Impact of Dataset Extractivity on Abstractive Summarization</i> Yubin Ge, Sullam Jeoung, Ly Dinh and Jana Diesner	13963
<i>Commonsense Knowledge Graph Completion Via Contrastive Pretraining and Node Clustering</i> Siwei Wu, Xiangqing Shen and Rui Xia	13977
<i>Incorporating Factuality Inference to Identify Document-level Event Factuality</i> Heng Zhang, Peifeng Li, Zhong Qian and Xiaoxu Zhu	13990
<i>Hybrid and Collaborative Passage Reranking</i> Zongmeng Zhang, Wengang Zhou, Jiaxin Shi and Houqiang Li	14003
<i>Sentence Embedding Leaks More Information than You Expect: Generative Embedding Inversion Attack to Recover the Whole Sentence</i> Haoran LI, Mingshi XU and Yangqiu Song	14022
<i>Learning Query Adaptive Anchor Representation for Inductive Relation Prediction</i> Zhiwen Xie, Yi Zhang, Jin Liu, Guangyou Zhou and Jimmy Xiangji Huang	14041
<i>Context or Knowledge is Not Always Necessary: A Contrastive Learning Framework for Emotion Recognition in Conversations</i> Geng Tu, Bin Liang, Ruibin Mao, Min Yang and Ruifeng Xu	14054
<i>Exploring Speaker-Related Information in Spoken Language Understanding for Better Speaker Diarization</i> Luyao Cheng, Siqi Zheng, Zhang Qinglin, Hui Wang, Yafeng Chen and Qian Chen	14068
<i>Cross-Lingual Knowledge Distillation for Answer Sentence Selection in Low-Resource Languages</i> Shivanshu Gupta, Yoshitomo Matsubara, Ankit R Chadha and Alessandro Moschitti	14078
<i>Run Like a Girl! Sport-Related Gender Bias in Language and Vision</i> Sophia Harrison, Eleonora Gualdoni and Gemma Boleda	14093
<i>People and Places of Historical Europe: Bootstrapping Annotation Pipeline and a New Corpus of Named Entities in Late Medieval Texts</i> Vit Novotny, Kristina Luger, Michal Štefánik, Tereza Vrabcová and Ales Horak	14104
<i>Check-COVID: Fact-Checking COVID-19 News Claims with Scientific Evidence</i> Gengyu Wang, Kate Rose Alderman Harwood, Lawrence G Chillrud, Amith Ananthram, Melanie Subbiah and Kathleen McKeown	14114
<i>Early Exit with Disentangled Representation and Equiangular Tight Frame</i> Yixin Ji, Jikai Wang, Juntao Li, Qiang Chen, Wenliang Chen and Min Zhang	14128
<i>Tokenization with Factorized Subword Encoding</i> David Samuel and Lilja Øvreliid	14143
<i>Rarely a problem? Language models exhibit inverse scaling in their predictions following few-type quantifiers</i> James Michaelov and Benjamin Bergen	14162
<i>A Little is Enough": Few-Shot Quality Estimation based Corpus Filtering improves Machine Translation</i> Akshay Batheja and Pushpak Bhattacharyya	14175
<i>How effective is machine translation on low-resource code-switching? A case study comparing human and automatic metrics</i> Li Nguyen, Christopher Bryant, Oliver Mayeux and Zheng Yuan	14186

<i>Images in Language Space: Exploring the Suitability of Large Language Models for Vision & Language Tasks</i>	
Sherzod Hakimov and David Schlangen	14196
<i>On the Expressivity Role of LayerNorm in Transformers' Attention</i>	
Shaked Brody, Uri Alon and Eran Yahav	14211
<i>DEnsity: Open-domain Dialogue Evaluation Metric using Density Estimation</i>	
ChaeHun Park, Seungil Chad Lee, Daniel Rim and Jaegul Choo	14222
<i>Fixing MoE Over-Fitting on Low-Resource Languages in Multilingual Machine Translation</i>	
Maha Elbayad, Anna Y Sun and Shruti Bhosale	14237
<i>Intent Discovery with Frame-guided Semantic Regularization and Augmentation</i>	
yajing sun, Rui Zhang, Jingyuan Yang and Wei Peng	14254
<i>An Empirical Comparison of LM-based Question and Answer Generation Methods</i>	
Asahi Ushio, Fernando Alva-Manchego and Jose Camacho-Collados	14262
<i>Contrastive Learning with Generated Representations for Inductive Knowledge Graph Embedding</i>	
Qian Li, Shafiq Joty, Daling Wang, Shi Feng, Yifei Zhang and Chengwei Qin	14273
<i>Decouple knowledge from paramters for plug-and-play language modeling</i>	
Xin Cheng, Yankai Lin, Xiuying Chen, Dongyan Zhao and Rui Yan	14288